arXiv:1606.05732v1 [cs.DS] 18 Jun 2016

How to Fake Multiply by a Gaussian Matrix

Michael Kapralov* EPFL
Vamsi K. Potluru† Comcast Cable
David P. Woodruff‡ IBM Research

June 21, 2016

Abstract

Have you ever wanted to multiply an $n \times d$ matrix X, with $n \gg d$, on the left by an $m \times n$ matrix $\tilde{G}$ of i.i.d. Gaussian random variables, but could not afford to do it because it was too slow? In this work we propose a new randomized $m \times n$ matrix T, for which one can compute $T \cdot X$ in only $O(\text{nnz}(X)) + \tilde{O}(m^{1.5} \cdot d^3)$ time, for which the total variation distance between the distributions $T \cdot X$ and $\tilde{G} \cdot X$ is as small as desired, i.e., less than any positive constant. Here $\text{nnz}(X)$ denotes the number of non-zero entries of X. Assuming $\text{nnz}(X) \gg m^{1.5} \cdot d^3$, this is a significant savings over the naïve $O(\text{nnz}(X)m)$ time to compute $\tilde{G} \cdot X$. Moreover, since the total variation distance is small, we can provably use $T \cdot X$ in place of $\tilde{G} \cdot X$ in any application and have the same guarantees as if we were using $\tilde{G} \cdot X$, up to a small positive constant in error probability. We apply this transform to nonnegative matrix factorization (NMF) and support vector machines (SVM).

1 Introduction

One approach to handle high dimensional data, often in the form of a matrix, is to first project the data to a much lower dimensional subspace. This is an example of sketching and the last decade has seen a systematic study of this approach. A linear sketch of a matrix replaces the original matrix by a smaller matrix which is often obtained by a random projection of the original matrix (see, e.g., Woodruff (2014) for a survey). Random projections have been successfully applied to speed up least squares regression and have been implemented with remarkable success Avron et al. (2010). This is impressive considering the fact that these solvers have been highly optimized over the last few decades, exploiting both algorithmic improvements and machine dependent optimizations.

Many of these works rely on *fast projection matrices*, such as the Subsampled Randomized Hadamard Transform or the CountSketch, the latter being particularly well-suited for sparse data (see, e.g., Woodruff (2014) and references therein). However, there are certain applications for which multiplying by a Gaussian matrix is the only way that is known to reduce the dimensionality of the data. This arises mainly because the application requires rotational symmetry, which is often not preserved by other fast transforms, or because additional properties, such as spreading out a sparse vector to a vector with non-spiky elements, do not hold for transforms like CountSketch (some of these hold for the Fast Hadamard Transform, but the latter are not known to be able to exploit sparsity). We give two such applications below, one to nonnegative matrix factorization (NMF), and one to support vector machines (SVM).

1.1 Our Results

A New Randomized Transform. In this work we propose a new randomized transform T, which we call the CountGauss. It is simply a product of a CountSketch matrix and a Gaussian matrix. That is, given an $n \times d$ matrix X which we would like to multiply by an $m \times n$ matrix $\tilde{G}$ of Gaussians, we instead let $T = G \cdot S$, where S is a $B \times d$ CountSketch matrix where $B = \tilde{O}(d^2 m^{\cdot 5})$, and $\tilde{G}$ is an $m \times B$ matrix of i.i.d. Gaussians. Recall that a CountSketch

*michael.kapralov@epfl.ch
†vamsi_potluru@cable.comcast.com
‡dpwoodru@us.ibm.com

matrix S satisfies that each column of S has only a single non-zero entry chosen in a uniformly random position. That non-zero is 1 with probability $1/2$, and -1 with probability $1/2$. The columns of S are independent of each other. Importantly, computing $S \cdot X$ can be done in $O(\text{nnz}(X))$ time, and this significantly reduces the number of rows of X. Then computing $G \cdot (S \cdot X)$ can now be done in $\tilde{O}(m^{1.5}d^3)$ time. While such a composition of matrices has been used before in the context of subspace embeddings for regression, see, e.g., Clarkson and Woodruff (2013b), here we show a new property of this composition - the distribution of $G \cdot S \cdot X$ looks like the distribution of $\tilde{G} \cdot X$! Formally, the statistical distance between the two distributions is smaller than any positive constant.

Therefore, in any application which uses $\tilde{G} \cdot X$, if we replace $\tilde{G} \cdot X$ with $G \cdot S \cdot X$, then if p is the success probability of the old algorithm, then the success probability of the new algorithm is at least $p - \delta$, where $\delta > 0$ is an arbitrarily small constant.

We now give applications.

Non-negative Matrix Factorization. Learning low rank structures and representations is a fundamental problem in machine learning. With the rise of data-driven decision making, many businesses, government agencies, and scientific laboratories are collecting increasingly large amounts of data each day. For instance, the large Hadron Collider (LHC) experiments represent about 150 million sensors acquiring around 40 million samples per second. Even working with 0.001 percent of the sensor data, the data flow from all four LHC experiments is around 25 petabytes per day Brumfiel (2011). This means the traditional approach of storing the data, and then processing it later, may be infeasible. One approach would be to subsample the incoming streams. However, we may lose valuable information in the form of infrequent events.

We use our transform to solve the nonnegative matrix factorization (NMF) problem. Previous approaches Damle and Sun (2014); Benson et al. (2014); Tepper and Sapiro (2015) have used random matrices for the projection. However, these approaches can be slow if the dimensionality of the data is high since they rely on multiplying by Gaussian matrices, e.g., for natural images or structural Magnetic Resonance Imaging brain scans. Recent work by Smola et al. Le et al. (2013) have shown that sometimes dense random Gaussian matrices can be replaced by faster transforms, and moreover, each row of the transform is equally likely to be in any direction on the unit sphere. To show the correctness of the NMF algorithm, however, we need a much stronger property than this, namely that any small subset of rows of the transform has the property that its product with a fixed matrix X has low variation distance to the distribution of a product of a Gaussian matrix with X. These latter properties, of having a fast transform with *equal representation of directions on the sphere*, do not seem to have been exploited in the context of NMF. Our transformation, since it has low variation distance to multiplying by a Gaussian matrix, directly applies here and we can use existing analysis.

We note that the classical way of speeding up Gaussian transforms via the Fast Hadamard or Fast Fourier Transform (see, e.g., Tropp (2011)) do not work in this context, since they miss large sections of the sphere, and we provide a formal counterexample in Section 9. Intuitively, while it is fine to miss directions along large sections of the sphere to approximate the norm of a vector, it is not fine to miss directions for NMF, where the corresponding polytope partitions the sphere into a small number of caps, and each cap should have a random direction chosen from it.

Support Vector Machines. We also apply random projections to the support vector machines (SVM) problem. Previously, the CountSketch (CW) Clarkson and Woodruff (2013a) projection and random Gaussian (RG) projection have been applied to the linear SVM problem. Despite Countsketch being much faster than the Gaussian projection, the overall running time of projection together with the SVM solver was similar for both projections Paul et al. (2014), since the training of the projected data was faster when using Gaussian projections. Our projection combines the CW matrix with a smaller Gaussian matrix thereby getting the best of both worlds — similar projection time as CountSketch and similar Gaussian properties of RG that are useful for SVM.

Experiments. We empirically validate our results for both NMF and SVM applications. For NMF, we give an experimental evaluation by comparing with state-of-the-art algorithms such as SPA Gillis et al. (2014), XRAY Kumar et al. (2013), naïve random projections Damle and Sun (2014) , structured Gaussian random projections Tepper and Sapiro (2015), and Tall-Skinny QR factorization Benson et al. (2014) for NMF problems with applications to breast cancer, flow cytometry, and climate data. Also, we show experimental speedups using our projection when combined with linear SVM solvers for document classification problems Paul et al. (2014).

2 A New Randomized Transform

A CountSketch matrix $S \in \mathbb{R}^{B \times n}$ is a matrix all of whose rows have exactly one nonzero in a uniformly random location, and the value of the nonzero element is independently chosen to be -1 or $+1$ with equal probability. We denote the number of rows in the CountSketch matrix by B.

We prove the next theorem, which gives the formal guarantees of our new transform.

Theorem 1. *There exists an absolute constant $C > 0$ such that for every $\delta \in (0,1)$, every integer $m \geq 1$ and every matrix $U \in \mathbb{R}^{n \times d}$ with orthonormal columns if $B \geq \frac{1}{\delta} C(\log n)^4 \cdot d^2 \cdot m^{1/2}$, $S \in \mathbb{R}^{B \times n}$ is a random CountSketch matrix, and $G \in \mathbb{R}^{m \times B}$ and $\tilde{G} \in \mathbb{R}^{m \times n}$ are matrices of i.i.d. unit variance Gaussians, then the total variation distance between the joint distribution GSU and $\tilde{G}U$ is less than δ.*

The proof is given in Section 7. We note that Theorem 1 applies to matrices U with orthonormal columns. This is sufficient for applying our transform to an arbitrary matrix X, since we can write $X = UR$, where the columns of U form an orthonormal basis for the range of X, and apply the theorem to U. Since $G \cdot S \cdot U$ is close the $\tilde{G} \cdot U$ in total variation distance, $G \cdot S \cdot UR = G \cdot S \cdot X$ is close to $\tilde{G} \cdot UR = \tilde{G} \cdot X$ in total variation distance as well. We note that the role of d and n in Theorem 1 is swapped in comparison to our notation for application to NMF below. The notation in Theorem 1 is more consistent with the numerical linear algebra literature, and we thus prefer to state the theorem in this form.

In Section 7 of our paper we also first present a weaker version of Theorem 1 (see Theorem 5) which establishes exactly the same guarantees but instead with $B \geq \frac{C}{\delta^2} \cdot d^2 \cdot m$. This version has a much simpler proof.

We first present the intuition behind the simpler version of Theorem 1. Consider the distribution of the first row of the two matrices, namely $\tilde{G}U$ versus GSU. Both random variables are Gaussians in dimension d, but while the former is an ideal isotropic Gaussian, the latter, despite being Gaussian, has correlated entries. The correlations between the entries are due to the fact that the CountSketch matrix S is not a perfect isometry: the correlation is given exactly by $U^T S^T SU$, which is the identity in expectation, but not for most realizations of S. In order to show that these two distributions are close in total variation distance, it would suffice to argue that the covariance matrix $U^T S^T SU$ is sufficiently close to the identity. This is exactly how the proof of the simpler version of Theorem 1 proceeds, which fixes an S for which $U^T S^T SU$ is sufficiently close to the identity, using a so-called "approximate matrix product" theorem in the linear algebra community. After fixing such an S, one can use that the rows of $G \cdot S$ and the rows of $\tilde{G}$ are independent, and then bound the variation distance between individual rows of $G \cdot S$ and of $\tilde{G}$. For the latter, it is convenient to work with Kullback-Leibler divergence (KL divergence) which is additive over product spaces; here we bound the KL divergence betwee a standard multivariate Gaussian and one with covariance matrix $U^T S^T SU$.

While this works for the simpler version of Theorem 1, it can be seen that since we need to ensure that the **joint** distribution of $\tilde{G}U$ is close to the joint distribution of GSU for most choices of S, we would need to set B to be at least $\approx d^2 m$ as opposed to $d^2\sqrt{m}$, as in our bound in Theorem 1. The idea behind the stronger result is to crucially use the fact that the distribution of GSU is a mixture of Gaussians with varying covariance matrices. When $B \approx d^2\sqrt{m}$, one can see that most Gaussians GSU in the mixture are too far in distribution from $\tilde{G}U$. However, we show that these differences that exist for **most** choices of the CountSketch matrix S cancel out **on average**. While mixtures of Gaussians with varying means and fixed covariance structure have been analyzed in the literature, to the best of our knowledge our analysis is the first to handle nontrivial mixtures with changing covariance.

3 Preliminaries for the Applications

A few applications of our new randomized transform are NMF and SVM, which we now formally define.

3.1 Nonnegative Matrix Factorization

Given a nonnegative matrix X of size $d \times n$, we would like to approximate it as a product of nonnegative matrices as follows: $X \approx WH$, where W is of size $d \times k$ and H is $k \times n$. This problem was studied by Paatero and Tapper Paatero and Tapper (1994) under the name of positive matrix factorization and gained a wider popularity through the work of Lee and Seung Lee and Seung (2001). NMF arises in a wide range of problems and application domains such as

curve resolution in chemometrics and document clustering; further references can be found in Arora et al. (2012). Various extensions to the original model to incorporate domain knowledge such as sparsity, orthogonality Ding et al. (2006), and under-approximation Gillis and Glineur (2010) have also been studied. Commonly used measures of approximation include the Frobenius norm, Itakuro-Saito (IS), and Bregman divergence with applications in image processing, speech and music analysis Yılmaz et al. (2011) among other places. Typical algorithms use alternating minimization to solve the non-convex objective function arising from NMF.

Until recently, the complexity of the NMF problem was unknown. Vavasis established that the NMF problem is NP-hard Vavasis (2009). However, if the data satisfies the separability condition, a condition introduced by Donoho and Stodden Donoho and Stodden (2003), then tractable algorithms exist and have been recently proposed by Arora et al. Arora et al. (2012); Recht et al. (2012). Formally, a nonnegative matrix X is k-separable if it satisfies the following condition: $X = X_I H$, where I is an index set of size k corresponding to the columns of the data matrix X. Geometrically, this assumption implies that the columns of X lie in a cone generated by the k selected columns of X indexed by I. One can view these k selected columns as the extreme points of a polytope containing all other columns. In practice, k is much smaller than both d and n. We will assume k-separability.

Given X_I, one can solve for X by solving a nonnegative least squares problem Damle and Sun (2014), and therefore our focus is on finding X_I, or equivalently, the index set I of extreme points of the point cloud formed by the columns of X.

To understand the guarantees of our algorithm, we define a few geometric notions also used in Damle and Sun (2014), which we refer to for more background. The *normal cone* of a convex set C at a point x is the cone

$$N_C(x) = \{w \in \mathbb{R}^d \mid w^T(y - x) \leq 0 \text{ for any } y \in C\},$$

that is, it is the cone defined by the outward normals of supporting hyperplanes at the point x. One can define a measure $\omega(K)$ on any cone K, which for full-dimensional cones K satisfies $\omega(K) = \Pr[\theta \in K \cap S^{d-1}]$ where θ is a uniformly random point on the sphere S^{d-1} in d dimensions. This measure is known as the *solid angle* of K. For any convex polytope C, if P is the set of its extreme points, then $\sum_{p \in P} \omega(N_C(p)) = 1$, that is, the solid angles of the normal cones at the extreme points sum to 1. If we label the points $p_i \in P$, we will use the shorthand $\omega_i = \omega(N_C(p_i))$.

A key property we will use is that for a unit vector u and a convex set C, the maximum inner product of u with any point $p \in C$ is achieved by an extreme point p of C. Moreover, the maximum is achieved by the extreme point p precisely when $u \in N_C(p)$. This follows since the inner product with a fixed vector u is a linear function, which is maximized by an extreme point for any convex set. These conditions also hold if we replace maximum with minimum.

Our results, as in Damle and Sun (2014), depend on the condition number $\kappa = \frac{1}{k \log\left(\frac{1}{\max_i 1-2\omega_i}\right)}$. The larger κ is, the more pointed the polytope defined by the columns of X is, whereas if κ is small, the polytope has "fatter" vertices.

3.2 Support Vector Machines

Given a dataset of samples and labels $\{x_i, y_i\}_{i=1}^N$ where x_i corresponds to sample i and y_i the corresponding label belonging to one of two classes denoted by $\{-1, 1\}$, we would like to find a maximum-margin hyperplane that separates the two classes. The primal form for the linear SVM problem is as follows:

$$\min_w \frac{1}{2}\|w\|_2^2 + \frac{C}{N}\sum_{i=1}^{N} \max(0, 1 - y_i\langle w, x_i\rangle) \tag{1}$$

where C is soft-margin parameter which allows for mis-classfication errors in the dataset and w is the maximum-margin hyperplane that we are learning from the data. The dual form for the linear SVM problem is given as follows:

$$\max_{0 \leq \alpha \leq C} 1^T\alpha - \frac{1}{2}\alpha^T Y X X^T Y \alpha \tag{2}$$

Previously Paul et al. (2014) have shown that the margin (hyperplane) and minimum enclosing ball of the original data is preserved after projection up to a multiplicative factor. However, in their original formulation it is possible to just replace all points with zero to achieve the same guarantee. We strengthen the theorems by requiring that the projected data upper bound the objective of the original data. The details are given in Section 10.

4 Application to NMF

We consider the separable NMF problem as defined in Section 3. We first review an algorithm proposed by Damle and Sun Damle and Sun (2014). Their algorithm involves the computation of $\tilde{G}X$ where $\tilde{G}$ is of dimensions $m \times d$ for a parameter m, and the element-wise entries are distributed as $N(0,1)$. Notice that we need to first compute the $m \times d$ random matrix $\tilde{G}$ which is itself dense. We also need to compute the matrix product $\tilde{G}X$ with the input data. This is computationally expensive and is of the order $O(mnd)$ in practice. Fast matrix multiplication routines Coppersmith and Winograd (1990); Williams (2012) can be used in theory, but the time will still be at least $k^{\omega-2}nd$, where $\omega \approx 2.376$ is the exponent of fast matrix multiplication. Instead, we propose to use our new transform to significantly speed up the computations for extracting the extreme points in the dataset. Note that both approaches are easily amenable to distributed-data settings by simply sharing the seed of the random number generator which allows identical matrix transformations on all the computational nodes. Our new algorithm is called **C**ount **G**auss NMF or **CountGauss** and is as follows: Instead of using Gaussian random matrices for the projection, we approximate them

Algorithm 1 **C**ount**G**auss NMF (**CG**)

Initialize the index sets I_{max}, I_{min} to empty.

1. Let $T = G \cdot S$ where G is an $m \times B$ matrix of i.i.d Gaussians, and S is a $B \times d$ CountSketch matrix. Here $B = Cn^2m^{\cdot 5}\log^4 n/\delta$.
2. Compute the product $Z = TX$.
3. Find the indices which give the maximum and minimum across each row of Z corresponding to I_{max}, I_{min}

by the following projection matrix $T = G \cdot S$, where the matrices are defined in Algorithm 1.

Consider the convex polytope defined by the columns of X and their negations. As defined in Section 3, we assume k-separability, namely, that there are k columns of X, indexed by I, for which $X = X_I H$ for a nonnegative matrix H. The columns of X_I are the extreme points of a convex polytope C. By definition of an extreme point of a convex polytope, the indices found in step 3 of Algorithm 1 belong to the index set I.

Damle and Sun show the following.

Theorem 2. *(Theorem 3.3 of Damle and Sun (2014)) Consider a modification to Algorithm 1 in which we replace T by an $m \times d$ matrix of i.i.d. $N(0,1)$ random variables, where $m = \kappa k \log(\frac{k}{\delta})$, where recall $\kappa = \frac{1}{k \max_i 1-2\omega_i}$ is the condition number. Then the probability that the output $I_{min} \cup I_{max}$ of Algorithm 1 contains the index set I of extreme points of X is at least $1 - \delta$.*

Using Theorem 1, we analyze the performance of Algorithm 1.

Theorem 3. *Let $\delta > 0$ be given. Suppose in Algorithm 1 we set the parameter $m = \kappa k \log(\frac{k}{\delta})$, where $\kappa = \frac{1}{k \max_i 1-2\omega_i}$ is the condition number, and choose $B \geq \frac{1}{\delta} C(\log d)^4 \cdot n^2 \cdot m^{1/2}$ for a sufficiently large constant $C > 1$ as per Theorem 1. Then the probability that the output $I_{min} \cup I_{max}$ of Algorithm 1 contains the index set I of extreme points of X is at least $1 - 2\delta$.*

Proof. Let I be the index set of extreme points of the polytope defined by the columns of X. By definition of an extreme point, in each iteration of step 4 of the algorithm, we add an index $i \in I$ to I_{max} and an index $j \in I$ to I_{min} (since we are taking the inner product with a linear function). Therefore, the behavior of Algorithm 1 is the same if we instead, in each invocation of step 2, compute the product $Z = TX_I$.

By our assumption on m, since X_I is a $d \times k$ matrix we may apply Theorem 1, with the role of n and d in that theorem swapped, to obtain that the variation distance of the distributions of TX_I and $\tilde{G}X_I$ is at most $1 - \delta$, where $\tilde{G}$ is a matrix of i.i.d. $N(0,1)$ random variables. Therefore, we can apply Theorem 2 to conclude by a union bound that the output of Algorithm 1 contains the set I with probability at least $1 - 2\delta$. □

We obtain the same guarantee as Theorem 2 with considerably faster computation time. Indeed, our matrix product TX can be computed in $O(\text{nnz}(X))+\tilde{O}(m^{1.5}n^3)$ time using our transform T, as opposed to the $O(dnm)$ time needed in Damle and Sun (2014) to compute the product $\tilde{G}X$ for a matrix $\tilde{G}$ of i.i.d. Gaussians. This is significant when d is very large.

Distributed Environments: Our results naturally provide solutions to NMF in a distributed environment in which the columns of X are partitioned across multiple servers. Indeed, the servers can agree upon a short random seed of length $O(d)$ words to generate T. Each server can then compute its local sets I_{max}, I_{min}, and send them to a coordinator who can find the global maxima and minima.

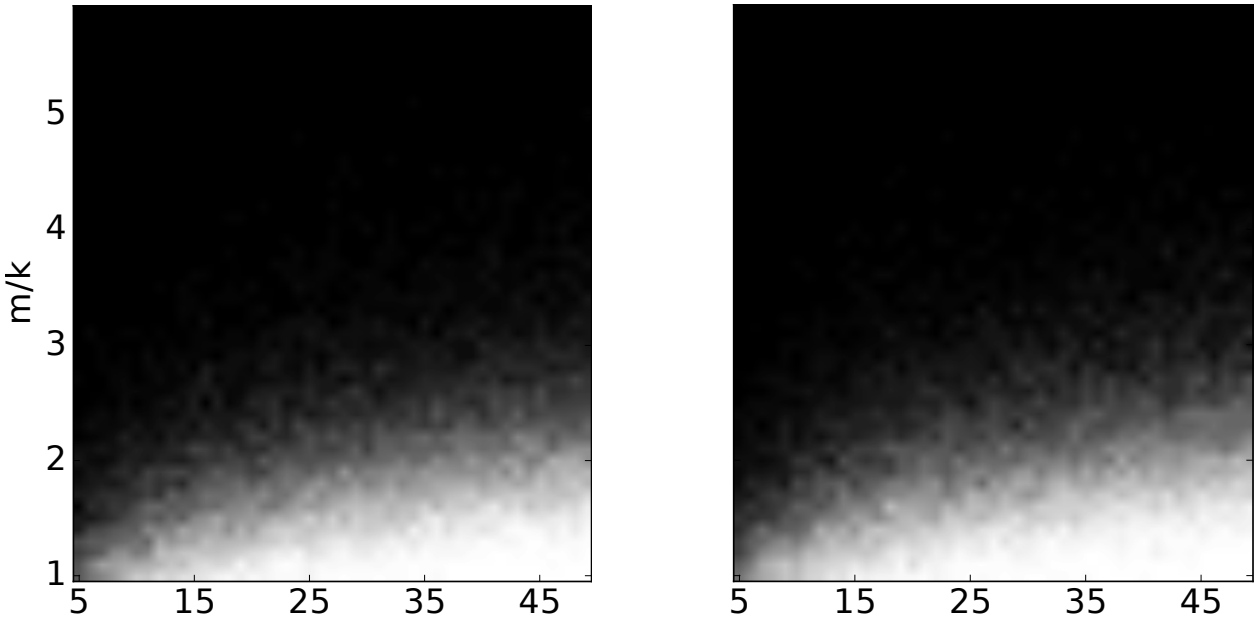

Figure 1: The fraction of trials in which the CG algorithm correctly extracted all "k' extreme points. For each value of k and m, we generated 500 matrices such that data matrix is of size 1000×500 and shown is how often we successfully recovered the original anchors (black indicates success). (Left) We contrast gaussian random projections (GP) with (right) our algorithm countGauss. Note that we recover the anchors with a similar success rate as GP.

5 Other Related work

Over the last couple of years, many approaches have been proposed to solve the separable-NMF problem.

XRAY Selects the anchors one at a time by expanding a cone until all columnds in the dataset are contained in it. At each step, XRAY finds the datapoint (column) which maximizes the inner product with the current residual matrix. It then computes the residual matrix corresponding to the new set of anchor points Kumar et al. (2013).

SPA Successive projection algorithm Araújo et al. (2001); Gillis et al. (2014) is a family of recursive algorithms where the projections are given by strongly convex functions.

TSQR Use tall and thin QR factorization when the number of rows/features is really large Benson et al. (2014). This approach is especially attractive when the number of features is really large ($\gg 10^6$) and number of samples is small ($< 10^5$).

SC In Tepper and Sapiro (2015), an algorithm similar to the one proposed by Damle and Sun Damle and Sun (2014) is proposed. The difference is that instead of choosing a Gaussian or FastFood projection matrix, the projection is chosen to be a matrix which depends on X (data dependent projection), namely, one that is found via the subspace power iteration (see Figure 3 of Tepper and Sapiro (2015)). This approach is expensive in the case of distributed settings since the projection matrix depends on all the samples.

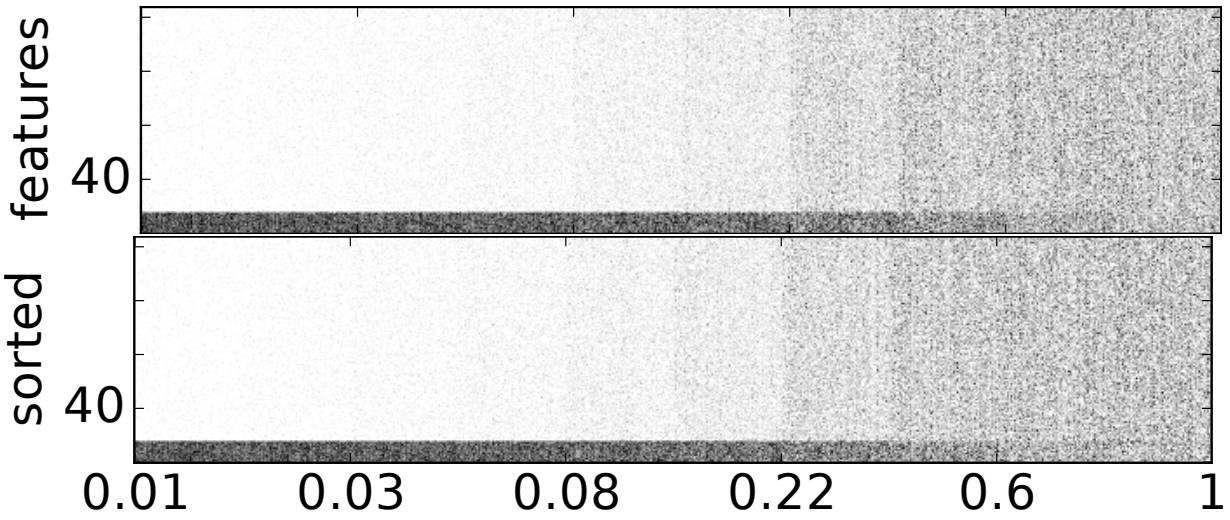

Figure 2: We show the scree plots at 20 noise levels and notice that there are sharp transitions at 20 corresponding to the rank of the data. (Top) Gaussian random projections and (bottom) our algorithm countGauss are applied to the dataset. For each noise value in $\{0.01, 0.02, 0.03, 0.05, 0.08, 0.12, 0.22, 0.36, 0.6, 1\}$, we generate 100 datasets. At higher noise levels, we note that both the algorithms GP and CG have most of the features active and there is no longer a sharp transition at 20.

Test		Proj		SVMf		Margin		Proj	Algo
mean	std	mean	std	mean	std	mean	std		
17.92	11.29	0.0000	0.0000	0.89	0.38	2.1057	3.9391	full	full
24.71	12.60	0.0086	0.0042	0.38	0.19	1.6792	3.5714	128	countSketch
25.27	13.08	0.0216	0.0047	**0.16**	**0.11**	1.6277	3.5634	128	countGauss
25.07	13.20	0.3676	0.1569	0.49	0.20	1.7143	3.7143	128	RG
17.92	11.29	0.0000	0.0000	0.89	0.38	2.1057	3.9391	full	full
22.56	12.42	0.0082	0.0036	0.54	0.21	1.8778	3.6709	256	countSketch
24.34	12.23	0.0565	0.0091	**0.21**	**0.07**	1.8722	3.7389	256	countGauss
23.66	12.86	0.8178	0.3286	0.98	0.35	1.8895	3.6747	256	RG
17.92	11.29	0.0000	0.0000	0.89	0.38	2.1057	3.9391	full	full
21.31	11.92	0.0075	0.0032	0.72	0.28	1.9914	3.7989	512	countSketch
22.11	12.89	0.1865	0.0228	**0.45**	**0.11**	1.9893	3.8453	512	countGauss
22.42	12.37	1.6057	0.6437	1.88	0.67	2.0148	3.9014	512	RG

Table 1: We applied CountGauss (CG), CountSketch (CW) and Random Gaussian (RG) on the TechTC300 dataset consisting of 295 pairs of data matrices and show the resulting mean and standard deviation for the resulting parameters such as projection time, SVMf time (projection + SVM training time), margin (gamma) and testing error. The results are shown over 10-fold cross validation with 4 repetitions and 3 runs over the random projection matrices. Note that the mean running times for our algorithm CG (highlighted) is faster than both CW and RG inspite of slower projection time than CW.

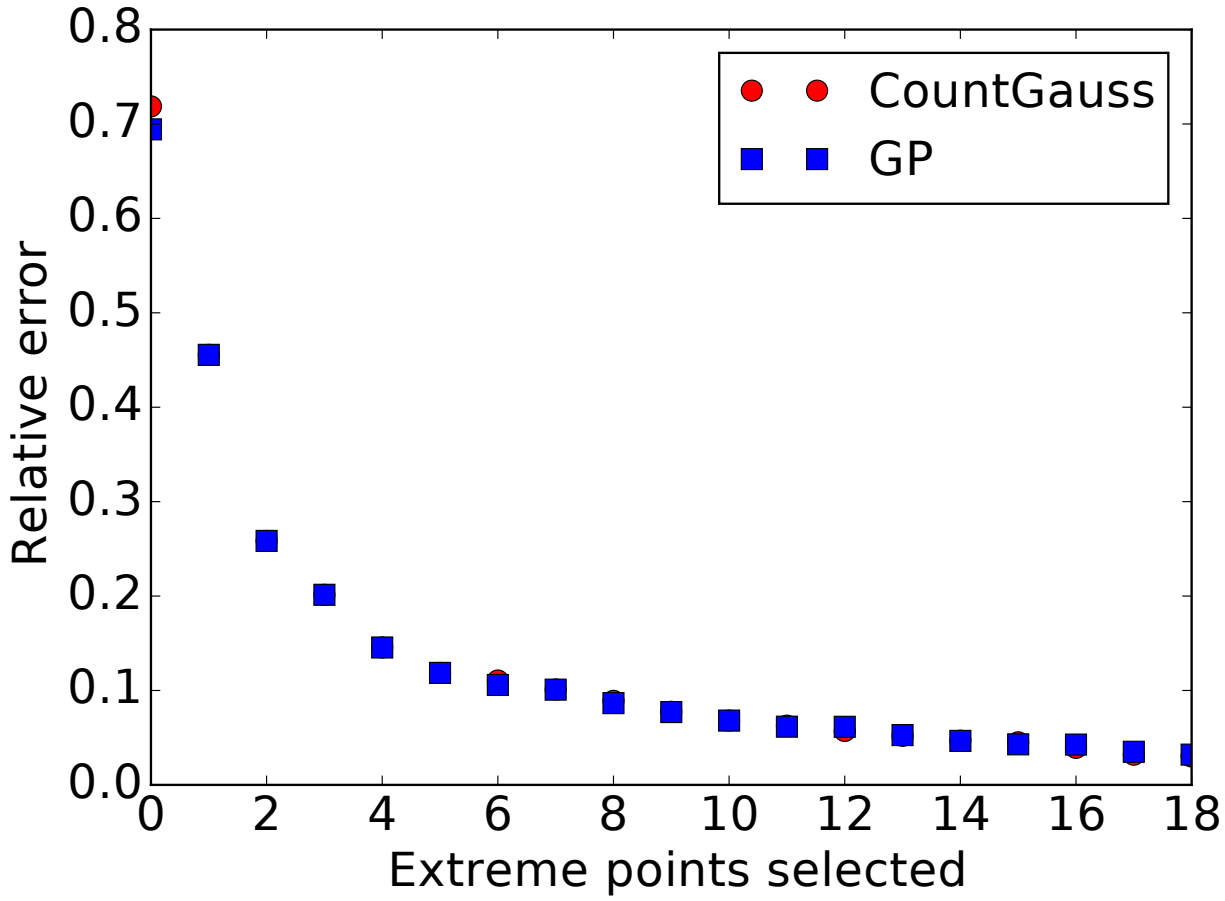

Figure 3: Relative reconstructive error as function of anchors selected by the two algorithms CountGauss and GP are shown. They are remarkably similar.

6 Experiments

We show experiments validating our projection operator countGauss (CG) for NMF problems on various synthetic and real-world datasets. Also, we apply CG on the SVM problem for the TechTC300 datasets. In all our experiments[1], we set $B = 5m$.

Synthetic datasets. Similar to Damle and Sun Damle and Sun (2014), we generate the data as follows: We set a grid of tuples (k, m) such that $m/k \approx \log k$. For each tuple, we generate 500 separable datasets, say X, such that they are of size 1000×500 and have nonnegative rank k. Set matrix U with i.i.d. samples from the uniform random distribution in $[0, 1]$ of size $d \times k$. Also, generate matrix V with the identity matrix for the top k indices and the rest with i.i.d samples from the uniform distribution. Normalize each row of the matrix V to unit norm and compute the matrix product $X = UV^\top$. From Figure 1, we see that the CountGauss algorithm also requires $O(k \log k)$ optimizations to find all k extreme points with high probability. We also test the algorithm in the noisy case. For that, we generate U of size 1000×20 with uniform entries in $[0, 1]$ and and set the first 20 columns of data matrix X to U. The remaining 190 columns of X are set to the midpoints of the $k(k+1)/2$-dimensional faces of the polytope with extreme points chosen by the first 20 columns of X. Now, we add Gaussian noise to X with noise level σ which creates many spurious extreme points. The resulting scree plot is shown in Figure 1.

Flow cytometry. The flow cytometry (FC) data represents abundances of fluorescent molecules labeling antibodies

[1] https://github.com/marinkaz/nimfa

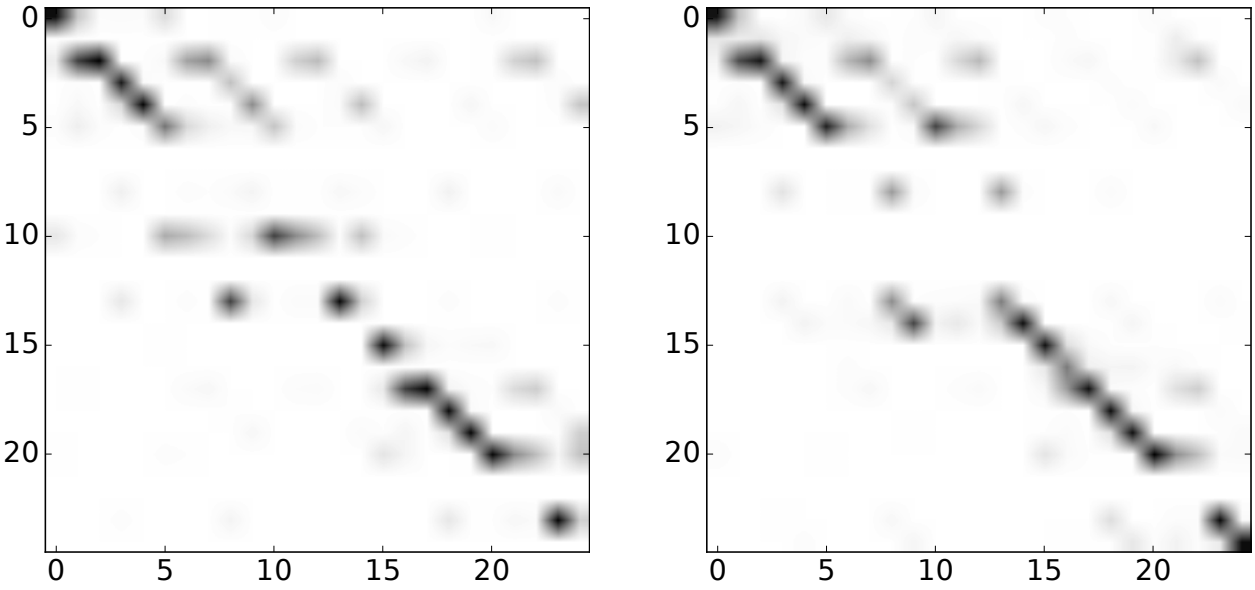

Figure 4: Coefficient matrices H are shown for the two algorithms GP and CountGauss for the flow cytometry data when k is set to 16. The coefficients tend to be clustered near the diagonal as has been previously observed.

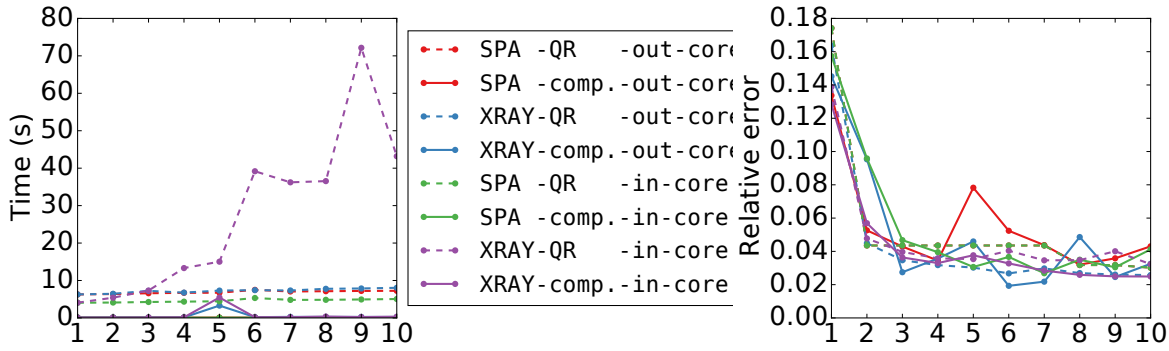

(A) Running times using CountGauss (comp==CG) (B) Relative error using CG

Figure 5: Extracting columns with CG versus using QR factorization. Note that the QR-based methods are optimized for tall-and-skinny matrices and tend to do poorly for fat matrices. Note that CG (and SC) tends to perform well since it is based on random projections and is at least an order of magnitude faster than QR-based methods.

that bind to specific targets on the surface of blood cells. A more detailed description of the dataset can be found in Benson et al. (2014). The measurements are represented as the data matrix A of size 40000×5. Since they study pairwise interactions in the data, the Kronecker product, $X = A \otimes A$ is formed which is of size $40000^2 \times 5^2$.

For this dataset, we exploit the data structure as follows. For some arbitrary input vector g, we know that $A \otimes Ag = A^\top GA$ where $g = \text{vec}(G)$. For each random projection, we can compute the matrix-matrix product AG very efficiently and in fact do not even need to generate the matrix G. For our algorithm, we do not need to explicitly compute the kronecker product and the complete NMF problem, including anchor selection and learning the weight coefficients, can be solved in a couple of seconds on an off-the-shelf desktop. As we can see from Figure 4 the results are pretty consistent from prior work Benson et al. (2014). The weight matrix H still maintains a diagonal-like structure as previously observed.

Gene expression breast cancer dataset. We utilize the hereditary breast cancer dataset collected by Hedenfalk et al.

(2001) which consists of the expression levels of 3226 genes on 22 samples from breast cancer patients. The patients consist of three groups: 7 patients with a BRCA1 mutation, 8 samples with a BRCA2 mutation4, and 7 additional patients with sporadic cancers. It was analyzed using separable NMF in Damle and Sun (2014) and we similarly preprocess the dataset by exponentiating to make the log- expression levels nonnegative and normalize the columns. The size of the data matrix is 3222×22. The result of applying our algorithm CG and GP are shown in Figure 3. Notice that we get similar reconstruction error as GP while we vary the number of anchors.

Climate Dataset. We obtained a climate dataset which was analyzed in Tepper and Sapiro (2015). The data size is 10512×23742. First we present the running times and reconstruction error when using SC versus QR-based algorithms and then show the corresponding results when using CG algorithm in Figure 5. Note that CG (and SC) which is based on random projections is an order-of-magnitude faster compared to QR factorization methods.

SVM TechTC-300 Dataset. We obtained the TechTC-300 dataset which is a comprehensive directory of the web. There are 295-pairs of categories which provides a rich framework for running SVM experiments Paul et al. (2014). Each data matrix has $10,000 - 40,000$ words and $150 - 280$ documents. LIBSVM was used with a linear kernel and soft-margin parameter C set to 500 for all experiments and we set the projections to $128, 256$, and 512. The results are summarized in the Table 1.

7 Proof of Theorem 1

The main result of this section is

Theorem 1 *(Restated) There exists an absolute constant $C > 0$ such that for every $\delta \in (0, 1)$, every integer $1 \leq m \leq n^4$ and every matrix $U \in \mathbb{R}^{n \times d}$ with orthonormal columns, if $B \geq \frac{1}{\delta} C(\log n)^4 \cdot d^2 \cdot m^{1/2}$, $S \in \mathbb{R}^{B \times n}$ is a random CountSketch matrix, and $G \in \mathbb{R}^{m \times B}$ and $\tilde{G} \in \mathbb{R}^{m \times n}$ are matrices of i.i.d. unit variance Gaussians, then the total variation distance between the joint distribution GSU and $\tilde{G}U$ is less than δ.*

Remark 4. *Note that we restrict the range of values of m in Theorem 1 to $[1 : n^4]$. This is because if $m > n^4$, the theorem requires $B \gg \frac{1}{\delta} n^2$, at which point the CountSketch matrix S becomes an isometry of $\mathbb{R}^n$ with high probability and the theorem follows immediately. At the same time restricting m to be bounded by a small polynomial of n simplifies the proof of Theorem 1 notationally.*

Recall that a CountSketch matrix $S \in \mathbb{R}^{B \times n}$ is a matrix all of whose columns have exactly one nonzero in a random location, and the value of the nonzero element is independently chosen to be -1 or $+1$. All random choices are made independently. Throughout this section we denote the number of rows in the CountSketch matrix by B. Note that the matrix S is a random variable. Let G denote an $m \times B$ matrix of independent Gaussians. For an $n \times d$ matrix U with orthonormal columns let $q : \mathbb{R}^d \to \mathbb{R}_+$ denote the p.d.f. of the random variable $G_1 SU$, where G_1 is the first row of G (all rows have the same distribution and are independent). We note that $G_1 SU$ is a mixture of Gaussians. Indeed, for any fixed S the distribution of $G_1 SU$ is normal with covariance matrix $(G_1 SU)^T (G_1 SU) = U^T S^T SU$. We denote the distribution of $G_1 SU$ given S by

$$q_S(x) := \frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2} x^T M^{-1} x}.$$

Throughout this section we use the notation $M := U^T S^T SU$. Note that since S is a random variable, M is as well. With this notation in place we have for any $x \in \mathbb{R}^d$

$$q(x) = \mathbf{E}_S \left[q_S(x) \right]. \tag{3}$$

Let $p : \mathbb{R}^d \to \mathbb{R}_+$ denote the pdf of the isotropic Gaussian distribution, i.e. for all $x \in \mathbb{R}^d$

$$p(x) = \frac{1}{\sqrt{(2\pi)^d}} e^{-\frac{1}{2} x^T x}. \tag{4}$$

Before giving a proof of Theorem 1, which is somewhat involved, we give a simple proof of a weaker version of the theorem, where the number of buckets B of our CountSketch matrix is required to be $\approx \frac{1}{\delta} d^2 m$ as opposed to $\approx \frac{1}{\delta} d^2 \sqrt{m}$:

Theorem 5. *There exists an absolute constant $C > 0$ such that for every $\delta \in (0,1)$, every integer $m \geq 1$ and every matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns if $B \geq \frac{1}{\delta^2} Cd^2 \cdot m$, $S \in \mathbb{R}^{B\times n}$ is a random CountSketch matrix, and $G \in \mathbb{R}^{m\times B}$ and $\tilde{G} \in \mathbb{R}^{m\times n}$ are matrices of i.i.d. unit variance Gaussians, then the total variation distance between the joint distribution GSU and $\tilde{G}U$ is less than δ.*

We will use the following measures of distance between two distribution in the proof of our main theorem (Theorem 1) as well as the proof of Theorem 5.

Definition 6 (Kullback-Leibler divergence)**.** *The Kullback-Leibler (KL) divergence between two random variables P, Q with probability density functions $p(x), q(x) \in \mathbb{R}^d$ is given by $D_{KL}(P||Q) = \int_{\mathbb{R}^d} p(x) \ln \frac{p(x)}{q(x)} dx$*

Definition 7 (Total variation distance)**.** *The total variation distance between two random variables P, Q with probability density functions $p(x), q(x) \in \mathbb{R}^d$ is given by $D_{TV}(P,Q) = \frac{1}{2}\int_{\mathbb{R}^d} |p(x) - q(x)|dx$.*

Theorem 8 (Pinsker's inequality)**.** *For any two random variables P, Q with probability density functions $p(x), q(x) \in \mathbb{R}^d$ one has $D_{TV}(P,Q) \leq \sqrt{\frac{1}{2} D_{KL}(P||Q)}$.*

The proof of Theorem 5 uses the following simple claim.

Claim 9 (KL divergence between multivariate Gaussians)**.** *Let $X \sim N(0, I_d)$ and $Y \sim N(0, \Sigma)$. Then $D_{KL}(X||Y) = \frac{1}{2}\mathrm{Tr}(\Sigma^{-1} - I) + \frac{1}{2}\ln\det\Sigma$.*

Proof. One has

$$\begin{aligned} D_{KL}(X||Y) &= \mathbf{E}_{X\sim N(0,I_d)}[-\frac{1}{2}X^TX + \frac{1}{2}X^T\Sigma^{-1}X + \frac{1}{2}\ln\det\Sigma] \\ &= \mathbf{E}_{X\sim N(0,I_d)}[\frac{1}{2}X^T(\Sigma^{-1} - I)X + \frac{1}{2}\ln\det\Sigma] \\ &= \frac{1}{2}\mathbf{E}_{X\sim N(0,I_d)}[\mathrm{Tr}((\Sigma^{-1} - I)XX^T)] + \frac{1}{2}\ln\det\Sigma \\ &= \frac{1}{2}\mathrm{Tr}(\Sigma^{-1} - I) + \frac{1}{2}\ln\det\Sigma, \end{aligned}$$

where we used the fact that for a vector X of independent Gaussians of unit variance one has $\mathbf{E}_X[X^TAX] = \mathrm{Tr}(A)$ for any symmetric A (by rotational invariance of the Gaussian distribution). □

We can now give

Proof of Theorem 5: One has by Lemma 23, **(1)** (see below; this is a standard property of the CountSketch matrix) that for any $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, and $B \geq 1$, if S is a random CountSketch matrix and $M = U^TS^TSU$, then $\mathbf{E}_S[||M - I||_F^2] = O(d^2/B)$. By Markov's inequality $\mathbf{Pr}_S[||I - M||_F > (2/\delta)\cdot O(d^2/B)] < \delta/2$. Let $\mathcal{E}$ denote the event that $||I - M||_F \leq (2/\delta)\cdot O(d^2/B)$. We condition on $\mathcal{E}$ in what follows. Since $B \geq \frac{1}{\delta^3} Cd^2m$ for a sufficiently large absolute constant $C > 1$, we have, conditioned on $\mathcal{E}$, that

$$||I - M||_F^2 \leq (2/\delta)\cdot O(d^2/B) = (2/\delta)\cdot\delta^3/(Cm) \leq 2\delta^2/(Cm). \tag{5}$$

Note that in particular we have $||I - M|| \leq ||I - M||_F < 1/2$ conditioned on $\mathcal{E}$ as long as $C > 1$ is larger than an absolute constant.

By Claim 9 we have $D_{KL}(X||Y) = \frac{1}{2}\mathrm{Tr}(I - \Sigma^{-1}) + \frac{1}{2}\ln\det\Sigma$. We now use Taylor expansions of matrix inverse and $\log\det$ provided by Claim 11 and Claim 12 (see below) to obtain

$$
\begin{aligned}
D_{KL}(X||Y) &= \frac{1}{2}\mathrm{Tr}(M^{-1}-I)+\frac{1}{2}\ln\det M \\
&= \frac{1}{2}\mathrm{Tr}\left(\sum_{k\geq 1}(I-M)^k\right)+\frac{1}{2}\sum_{k\geq 1}\left(-\mathrm{Tr}((I-M)^k)/k\right) \\
&= \frac{1}{2}\mathrm{Tr}\left(\sum_{k\geq 2}(I-M)^k\right)+\frac{1}{2}\sum_{k\geq 2}\left(-\mathrm{Tr}((I-M)^k)/k\right) \\
&= O(\mathrm{Tr}((I-M)^2)) \qquad \text{(since } ||I-M||_2 \leq ||I-M||_F < 1/2) \\
&= O(||I-M||_F^2) \\
&= O(2\delta^2/(Cm)) \qquad \text{(by (5))} \\
&\leq (\delta/4)^2/m \qquad (6)
\end{aligned}
$$

as long as $C > 1$ is larger than an absolute constant. This shows that for every $S \in \mathcal{E}$ one has $D_{KL}(p||q_S) \leq (\delta/4)^2/m$, and thus $D_{KL}(p||\tilde{q}|\mathcal{E}]) \leq (\delta/4)^2/m$, where we let $\tilde{q}(x) := \mathbf{E}_S[q_S(x)|\mathcal{E}]$.

We now observe that the vectors $(G_iSU)_{i=1}^m$ and $(\tilde{G}_iU)_{i=1}^m$ are vectors of independent samples from distributions $q(x)$ and $p(x)$ respectively. We denote the corresponding product distributions by q^m and p^m. Since the good event $\mathcal{E}$ constructed above occurs with probability at least $1-\delta/2$, it suffices to consider the distributions $\tilde{q}(x)$ and $p(x)$, as

$$D_{TV}(q^m,p^m) \leq \mathbf{Pr}[\bar{\mathcal{E}}] + D_{TV}(q^m,p^m|\mathcal{E}) = \mathbf{Pr}[\bar{\mathcal{E}}] + D_{TV}(\tilde{q}^m,p^m),$$

where $D_{TV}(q^m,p^m|\mathcal{E}) = D_{TV}(\tilde{q}^m,p^m)$ stands for the total variation distance between the distribution of $(\tilde{G}_iU)_{i=1}^m$ and the distribution of $(G_iSU)_{i=1}^m$ conditioned on $S \in \mathcal{E}$. We can now use the estimate from (6) to get

$$
\begin{aligned}
D_{TV}(\tilde{q}^m,p^m) &\leq \sqrt{\frac{1}{2}D_{KL}(p^m||\tilde{q}^m)} \quad \text{(by Pinsker's inequality)} \\
&= \sqrt{\frac{m}{2}D_{KL}(p||\tilde{q})} \quad \text{(by additivity of KL divergence over product spaces)} \\
&\leq \sqrt{\frac{m}{2}\cdot(\delta/4)^2/m} \quad \text{(by (6))} \\
&\leq \delta/4.
\end{aligned}
$$

□

The main source of hardness in proving the stronger result provided by Theorem 1 comes from the fact that unlike the setting of Theorem 5, where most elements in the mixture are close to isotropic Gaussians in KL divergence, in the setting of Theorem 1 most elements of the mixture are too far from isotropic Gaussians to establish our result directly (this can be seen by verifying that the bounds of Theorem 5 on the KL divergence of q_S to p are essentially tight). Thus, the main technical challenge in proving Theorem 1 consists of analyzing the effect of averaging over random CountSketch matrices that is involved in the definition of $q(x)$ in (3). The core technical result behind the proof of Theorem 1 is Lemma 10, stated below. Ideally, we would like a lemma that states that the ratio of the pdfs $q(x)/p(x)$ is very close to 1 for 'typical' values of x (for appropriate definition of a set of 'typical' x). Unfortunately, it is not clear how to achieve this result for the distribution $q(x)$ defined in (3). The problem is that some choices of CountSketch matrices S may lead to degenerate Gaussian distributions that are hard to analyze. For example, when S is not a subspace embedding, the matrix M may even be rank-deficient, and the inverse M^{-1} is then ill-defined. To avoid these issues, we work with an alternative definition. Specifically, instead of averaging the distributions $\frac{1}{\sqrt{(2\pi)^d \det M}}e^{-\frac{1}{2}x^TM^{-1}x}$ over all CountSketch matrices, we define a high probability event $\mathcal{E}$ in the space of matrices S (see Lemma 10 for the definition) and reason about the modified distribution $\tilde{q}(x)$ defined as

$$\tilde{q}(x) = \mathbf{E}_S\left[\frac{1}{\sqrt{(2\pi)^d \det M}}e^{-\frac{1}{2}x^TM^{-1}x}\middle|\mathcal{E}\right]. \qquad (7)$$

For technical reasons it turns out to be useful to define yet another distribution

$$q'(x) = \mathbf{E}_S \left[\frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2}x^T M^{-1} x} \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E} \right] + \xi \cdot p(x), \tag{8}$$

where $\xi = \mathbf{E}_S \left[\mathbf{Pr}_{X \sim q_S}[X \notin \mathcal{T}(S,U)] | \mathcal{E} \right] \leq n^{-20}$ and for each $S \in \mathcal{E}$ and U with orthonormal columns the set $\mathcal{T}(S,U)$ (see Definition 14) is an appropriately defined set of $x \in \mathbb{R}^d$ that are 'typical' for S and U. We first note that q' is indeed the p.d.f. of a distribution. First, it is clear that $q'(x) \geq 0$ for all x. Second, we have

$$\begin{aligned}
\int_{\mathbb{R}^d} q'(x)dx &= \int_{\mathbb{R}^d} \mathbf{E}_S \left[\frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2}x^T M^{-1} x} \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E} \right] + \xi \cdot \int_{\mathbb{R}^d} p(x)dx \\
&= 1 - \int_{\mathbb{R}^d} \mathbf{E}_S \left[\frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2}x^T M^{-1} x} \cdot \mathbf{I}[x \notin \mathcal{T}(S,U)] \middle| \mathcal{E} \right] + \xi \\
&= 1 - \mathbf{E}_S \left[\int_{\mathbb{R}^d} \frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2}x^T M^{-1} x} \cdot \mathbf{I}[x \notin \mathcal{T}(S,U)] dx \middle| \mathcal{E} \right] + \xi \\
&= 1 - \mathbf{E}_S \left[\int_{\mathbb{R}^d \setminus \mathcal{T}(S,U)} \frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2}x^T M^{-1} x} dx \middle| \mathcal{E} \right] + \xi \\
&= 1 - \mathbf{E}_S \left[\mathbf{Pr}_{X \sim q_S}[X \notin \mathcal{T}(S,U)] | \mathcal{E} \right] + \xi \\
&= 1, \quad \text{(by definition of } \xi)
\end{aligned}$$

as required.

As we show below, the total variation distance between q' and $\tilde{q}$ is a small n^{-10}, so working with q' suffices. The main argument of our proof shows that the distribution $q'(x)$ is close to $p(x)$ for 'typical' $x \in \mathbb{R}^d$. Then since q' is close to $\tilde{q}$ and the event $\mathcal{E}$ occurs with high probability, this suffices for a proof of Theorem 1. Formally, the core technical result behind the proof of Theorem 1 is

Lemma 10. *There exists an absolute constant $C > 0$ such that for every $\delta \in (0,1)$ and every matrix $U \in \mathbb{R}^{n \times d}$ with orthonormal columns if $B \geq \frac{1}{\delta} C (\log n)^4 d^2$ there exists a set $\mathcal{E}$ of CountSketch matrices and a subset $\mathcal{T}^* \subseteq \mathbb{R}^d$ that satisfies $\mathbf{Pr}_{X \sim p}[X \notin \mathcal{T}^*] \leq n^{-10}$ and $\mathbf{Pr}_{X \sim \tilde{q}}[X \notin \mathcal{T}^*] \leq n^{-10}$ such that if $S \in \mathbb{R}^{B \times n}$ is a random CountSketch matrix, then* **(1)** $\mathbf{Pr}_S[\mathcal{E}] \geq 1 - \delta/3$*, and* **(2)** *for all $x \in \mathcal{T}^*$ one has*

$$\left| \frac{q'(x)}{p(x)} - 1 \right| \leq O((d^2 \log^4 n)/B) + O(n^{-10}).$$

We now prove Theorem 1 assuming Lemma 10 and Claim 17. After this, we then prove Lemma 10 and Claim 17. We now give

Proof of Theorem 1: The proof relies on the observation that the vectors $(G_i S U)_{i=1}^m$ and $(\tilde{G}_i U)_{i=1}^m$ are vectors of independent samples from distributions $q(x)$ and $p(x)$ respectively. We denote the corresponding product distributions by q^m and p^m. Since the good event $\mathcal{E}$ constructed in Lemma 10 occurs with probability at least $1 - \delta/3$, it suffices to consider the distributions $\tilde{q}(x)$ and $p(x)$, as

$$D_{TV}(q^m, p^m) \leq \mathbf{Pr}[\bar{\mathcal{E}}] + D_{TV}(\tilde{q}^m, p^m | \mathcal{E}), \tag{9}$$

where $D_{TV}(\tilde{q}^m, p^m | \mathcal{E})$ stands for the total variation distance between the distribution of $(\tilde{G}_i U)_{i=1}^m$ and the distribution of $(G_i S U)_{i=1}^m$ conditioned on $S \in \mathcal{E}$. Further, we have by the triangle inequality

$$D_{TV}(\tilde{q}^m, p^m | \mathcal{E}) \leq D_{TV}((q')^m, p^m | \mathcal{E}) + D_{TV}(\tilde{q}^m, (q')^m | \mathcal{E}) \leq D_{TV}((q')^m, p^m | \mathcal{E}) + m \cdot n^{-10}, \tag{10}$$

since $D_{TV}(\tilde{q}^m, (q')^m | \mathcal{E}) \leq m D_{TV}(\tilde{q}, q' | \mathcal{E}) \leq m n^{-10}$, where $D_{TV}(\tilde{q}, q' | \mathcal{E}) \leq n^{-10}$ by Claim 17 below.

We first prove, using Lemma 10, that the KL divergence between $p(x)$ and $q'(x)$ restricted to the set $\mathcal{T}^*$ (whose existence is guaranteed by Lemma 10) is bounded by $O(((d\log n)^2/B)^2)$. Specifically, let

$$p_*(x) := \begin{cases} p(x)/\mathbf{Pr}_{X\sim p}[\mathcal{T}^*] & \text{if } x \in \mathcal{T}^* \\ 0 & \text{o.w.} \end{cases} \tag{11}$$

and

$$q'_*(x) := \begin{cases} q'(x)/\mathbf{Pr}_{X\sim q'}[\mathcal{T}^*] & \text{if } x \in \mathcal{T}^* \\ 0 & \text{o.w.} \end{cases} \tag{12}$$

Since $\mathcal{T}^*$ occurs with probability at least $1 - 1/n^{10}$ under both $\tilde{q}(x)$ and $p(x)$ by Lemma 21, it suffices to bound the total variation distance between the product of m independent copies of $q'_*(x)$ and m independent copies of $p_*(x)$. Specifically,

$$\begin{aligned} D_{TV}((q')^m, p^m|\mathcal{E}) &\le D_{TV}((q'_*)^m, p_*^m|(\mathcal{T}^*)^m) + m\mathbf{Pr}[q'(\mathbf{R}^d \setminus \mathcal{T}^*)] + m\mathbf{Pr}[p(\mathbb{R}^d \setminus \mathcal{T}^*)] \\ &\le D_{TV}((q'_*)^m, p_*^m) + 2mn^{-10}, \quad \text{(by Lemma 21)} \end{aligned} \tag{13}$$

where we used the fact that q'_* and p_* are supported on $\mathcal{T}^*$. Note that both distributions are still product distributions. By Pinkser's inequality and the product structure we thus get

$$\begin{aligned} D_{TV}((q'_*)^m, p_*^m) &\le \sqrt{\frac{1}{2} D_{KL}((q'_*)^m||p_*^m)} \quad \text{(by Pinsker's inequality)} \\ &= \sqrt{\frac{m}{2} D_{KL}(q'_*||p_*)} \quad \text{(by additivity of KL divergence over product spaces)} \end{aligned} \tag{14}$$

In what follows we bound $D_{KL}(q'_*||p_*)$. By Lemma 10 we have for every $x \in \mathcal{T}^*$ that

$$|q'(x)/p(x) - 1| \le O((d^2\log^4 n)/B) + O(n^{-10}), \tag{15}$$

so

$$\begin{aligned} |q'_*(x)/p_*(x) - 1| &= \left|(q'(x)/p(x)) \cdot \frac{\mathbf{Pr}_{X\sim q'}[\mathcal{T}^*]}{\mathbf{Pr}_{X\sim p}[\mathcal{T}^*]} - 1\right| = \frac{\mathbf{Pr}_{X\sim q'}[\mathcal{T}^*]}{\mathbf{Pr}_{X\sim p}[\mathcal{T}^*]} \cdot \left|(q'(x)/p(x)) - \frac{\mathbf{Pr}_{X\sim p}[\mathcal{T}^*]}{\mathbf{Pr}_{X\sim q'}[\mathcal{T}^*]}\right| \\ &\le \frac{\mathbf{Pr}_{X\sim q'}[\mathcal{T}^*]}{\mathbf{Pr}_{X\sim p}[\mathcal{T}^*]} \cdot \left(|q'(x)/p(x) - 1| + \left|1 - \frac{\mathbf{Pr}_{X\sim p}[\mathcal{T}^*]}{\mathbf{Pr}_{X\sim q'}[\mathcal{T}^*]}\right|\right) \\ &= (1 + O(n^{-10})) \cdot \left(|q'(x)/p(x) - 1| + O(n^{-10})\right) \\ &= O((d^2\log^4 n)/B) + O(n^{-10}). \quad \text{(by (15))} \end{aligned}$$

Since $B \ge \frac{1}{\delta} C d^2 \log^4 n$ for a sufficiently large constant $C > 0$ by assumption of the theorem, we get that

$$O((d^2\log^4 n)/B) + O(n^{-10}) < O(1/C) + O(n^{-10}) < 1/2.$$

We thus get, using the bound $|1/(1+x) - 1| \le 2|x|$ for $|x| \le 1/2$,

$$\begin{aligned} |p_*(x)/q'_*(x) - 1| &= \left|\frac{1}{q'_*(x)/p_*(x)} - 1\right| = \left|\frac{1}{1 + (q'_*(x)/p_*(x) - 1)} - 1\right| \\ &= O\left(|q'_*(x)/p_*(x) - 1|\right) \\ &= O((d^2\log^4 n)/B) + O(n^{-10}) \end{aligned} \tag{16}$$

We now use the fact that $|\ln(1+x) - x| \le 2x^2$ for all $x \in (-1/10, 1/10)$ to upper bound $D_{KL}(q'_*||p_*)$. Specifically,

we have

$$\begin{aligned} D_{KL}(q'_*||p_*) = \mathbf{E}_{X\sim q'_*}[\ln(q'_*(X)/p_*(X))] &\leq -\mathbf{E}_{X\sim q'_*}[\ln(p_*(X)/q'_*(X))] \\ &\leq -\mathbf{E}_{X\sim q'_*}[(p_*(x)/q'_*(x)-1)-(p_*(x)/q'_*(x)-1)^2] \\ &\leq -\mathbf{E}_{X\sim q'_*}[p_*(x)/q'_*(x)-1]+\mathbf{E}_{X\sim q'_*}[(p_*(x)/q'_*(x)-1)^2] \\ &= -(1-1)+\mathbf{E}_{X\sim q'_*}[(p_*(x)/q'_*(x)-1)^2] \\ &= \mathbf{E}_{X\sim q'_*}[(p_*(x)/q'_*(x)-1)^2] \\ &= O(((d^2\log^4 n)/B)^2+n^{-10}) \qquad \text{(by (16))} \end{aligned} \tag{17}$$

Since $B \geq \frac{1}{\delta}C(\log n)^4 d^2 \cdot m^{1/2}$ for a sufficiently large constant $C > 0$ by assumption of the theorem, substituting the bound of (17) into (14), we get

$$D_{TV}((q'_*)^m, p_*^m) \leq \sqrt{\frac{m}{2}D_{KL}(q'_*||p_*)} \leq \sqrt{\frac{m}{2}\cdot O(((d^2\log^4 n)/B)^2+n^{-10})} \leq \sqrt{\frac{m}{2}\cdot\delta^2/(8m)} \leq \delta/2.$$

Putting this together with (13), (10) and (9) using the assumption that $m \leq n^4$ gives the result. □

The rest of the section is devoted to proving Lemma 10, i.e. bounding

$$q'(x)/p(x) = \mathbf{E}_S\left[\exp\left(\frac{1}{2}x^Tx - \frac{1}{2}x^TM^{-1}x - \frac{1}{2}\log\det M\right)\cdot\mathbf{I}[x\in\mathcal{T}(S,U)]\middle|\,\mathcal{E}\right]+\xi, \tag{18}$$

where $\xi = \mathbf{E}_S\left[\mathbf{Pr}_{X\sim q_S}[X\notin\mathcal{T}(S,U)]\,|\,\mathcal{E}\right] \leq n^{-20}$, for 'typical' x sampled from the Gaussian distribution (i.e. $x \in \mathcal{T}^*$ – see formal definition below).

Organization. The rest of this section is organized as follows. We start by defining the set $\mathcal{E}$ of 'nice' CountSketch matrices in section 7.1, and proving that a random CountSketch matrix is likely to be 'nice'. We will in fact define a parameterized set $\mathcal{E}(\gamma)$ in terms of a parameter γ. In section 7.2 we define, for each matrix U (which can be thought of as fixed throughout our analysis) with orthonormal columns and CountSketch matrix S, a set $\mathcal{T}(S,U)$ of $x \in \mathbb{R}^d$ that are 'typical' for S and U. The ratio of pdfs in (18) can be approximated well by a Taylor expansion **for such 'typical'** $x \in \mathcal{T}(S,U)$. These Taylor expansions are developed in section 7.3 and form the basis of our proof. Unfortunately, these Taylor expansions are valid only for $x \in \mathcal{T}(S,U)$, i.e. for x that are 'typical' with respect to a given S. To complete the proof, we need to construct a universal 'typical' set $\mathcal{T}^*(U,\gamma)$ of $x \in \mathbb{R}^d$, again parameterized in terms of a parameter γ, that will allow for approximation via Taylor expansions for **all** $x \in \mathcal{T}^*(U,\gamma)$ **and** $S \in \mathcal{E}(\gamma)$. We construct such a set $\mathcal{T}^*(U,\gamma)$ in section 7.4. Finally, the proof of Lemma 10 is given in section 7.5.

7.1 Typical set $\mathcal{E}$ of CountSketch matrices and its properties

Our analysis of (18) starts by Taylor expanding M^{-1} and $\det M$ around the identity matrix. We now state the Taylor expansions, and the define a (family of) high probability events $\mathcal{E}(\gamma)$ (equivalently, sets of 'typical' CountSketch matrices) such that the Taylor expansions are valid for matrices $M \in \mathcal{E}(\gamma)$ for all sufficiently small γ.[2] The Taylor expansions that we use are given by

Claim 11. *For any matrix M with $||I-M|| < 1/2$ one has $M^{-1} = (I-(I-M))^{-1} = \sum_{k\geq 0}(I-M)^k$.*

Claim 12. *For any matrix M with $||I-M|| < 1/2$ one has $\log\det M = \log\det(I-(I-M)) = \sum_{k\geq 1} -\mathrm{Tr}((I-M)^k)/k$.*

For a parameter $\gamma \in (0,1)$ that we will later set to $1/\mathrm{poly}(\log n)$, define event $\mathcal{E}(\gamma)$ as

$$\mathcal{E}(\gamma) := \left\{||I-M||_F^2 \leq \gamma^2 \quad \text{and} \quad |\mathrm{Tr}(I-M)| \leq \gamma\right\}. \tag{19}$$

The events $\mathcal{E}(\gamma)$ occur with high probability even for fairly small γ as long as B is sufficiently large:

[2]Note that we use the notation $S \in \mathcal{E}(\gamma)$ and $M \in \mathcal{E}(\gamma)$ interchangeably. This is fine since $M = U^TS^TSU$ and the matrix U is fixed.

Claim 13. *For any matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, any $B \times n$ CountSketch matrix S we have* $\mathbf{Pr}[\mathcal{E}(\gamma)] \geq 1 - 3(d/\gamma)^2/B$.

Proof. By Lemma 23 below, we have $\mathbf{E}_S[||I-M||_F^2] \leq 2d^2/B$. Applying Markov's inequality to $||I-M||_F^2$, we get

$$\mathbf{Pr}[||I-M||_F^2 \geq \gamma^2] \leq \mathbf{Pr}[||I-M||_F^2 \geq \gamma^2(B/(2d^2)) \cdot \mathbf{E}[||I-M||_F^2]] \leq 2(d/\gamma)^2/B$$

as required.

We also have by Lemma 23 (fifth bound) that $\mathbf{E}_S[(\mathrm{Tr}(I-M))^2] \leq d^2/B$. Applying Markov's inequality to $(\mathrm{Tr}(I-M))^2$, we get

$$\mathbf{Pr}[|\mathrm{Tr}(I-M)| \geq \gamma] = \mathbf{Pr}[(\mathrm{Tr}(I-M))^2 \geq \gamma^2] \leq \mathbf{Pr}[(\mathrm{Tr}(I-M))^2 \geq \gamma^2(B/(d^2))\cdot\mathbf{E}[(\mathrm{Tr}(I-M))^2]] \leq (d/\gamma)^2/B.$$

A union bound over the two events gives the result. □

7.2 Typical sets $\mathcal{T}(S,U)$ and their properties

In order to construct a single typical set $\mathcal{T}^*$, we will need the following simple definitions of sets $\mathcal{T}(S,U)$ of $x \in \mathbb{R}^d$ that are 'typical' for a given CountSketch matrix (as opposed to the set $\mathcal{T}^*$ whose existence is guaranteed by Lemma 10, which contains x that are 'typical' for **all matrices** $S \in \mathcal{E}$ **simultaneously**). We will use

Definition 14 (Typical x)**.** *For any orthonormal matrix $U \in \mathbb{R}^{n\times d}$ and CountSketch matrix S we define*

$$\mathcal{T}(S,U) := \left\{ x \in \mathbb{R}^d : |x^T(I-M)x| \leq \frac{1}{100} \text{ and } |x^T(I-M)^2x| \leq \frac{1}{100} \right\},$$

The following claim will be useful in what follows. Its (simple) proof is given in Section 8:

Claim 15. *For any matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns and any CountSketch matrix $S \in \mathbb{R}^{B\times n}$ one has $||I-M||_F^2 \leq 4n^3$.*

The following claim is crucial to our analysis. A detailed proof is given in Section 8.

Claim 16. *For any matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, every $\gamma \leq 1/\log^2 n$, every CountSketch matrix $S \in \mathcal{E}(\gamma)$ one has* **(1)** $\mathbf{Pr}_{X\sim N(0,I_d)}[X \notin \mathcal{T}(S,U)] < n^{-40}$ *and* **(2)** *for any CountSketch matrix $S' \in \mathcal{E}(\gamma)$, $M' = U^TS'^TS'U$ one has $\mathbf{Pr}_{X\sim N(0,M')}[X \notin \mathcal{T}(S,U)] < n^{-40}$ for sufficiently large n.*

Using the claim above we get

Claim 17. *The total variation distance between $\tilde{q}$ (defined in* (7)*) and q' (defined in* (8)*) is at most n^{-10}. Further, $\xi \leq n^{-40}$.*

Proof. We have

$$\begin{aligned}
D_{TV}(\tilde{q},q') \leq 2\xi &\leq 2\int_{\mathbb{R}^d} \mathbf{E}_S\left[\frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2}x^TM^{-1}x} \cdot \mathbf{I}[x \notin \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] dx \\
&= 2\mathbf{E}_S\left[\int_{\mathbb{R}^d} \frac{1}{\sqrt{(2\pi)^d \det M}} e^{-\frac{1}{2}x^TM^{-1}x} \cdot \mathbf{I}[x \notin \mathcal{T}(S,U)]dx \middle| \mathcal{E}(\gamma)\right] \\
&= 2\mathbf{E}_S\left[\mathbf{Pr}_{X\sim N(0,M)}[x \notin \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] \\
&\leq 2n^{-40} \leq n^{-10} \qquad \text{(by Claim 16)}
\end{aligned}$$

as required. □

7.3 Basic Taylor expansions

In this section we define the basic Taylor expansions of $\tilde{q}(x)/p(x)$ that form the foundation of our analysis. Our analysis of (18) proceeds by first Taylor expanding M^{-1} and $\det M$ around the identity matrix using Claims 11 and 12, which is valid since for any $S \in \mathcal{E}(\gamma)$ for $\gamma < 1/2$ one has $||I - M||_2 \leq ||I - M||_F \leq 1/2$. This gives

$$\begin{aligned}
\tilde{q}(x)/p(x) &= \mathbf{E}_S \left[\exp\left(\frac{1}{2}x^T x - \frac{1}{2}\left(\sum_{k\geq 0} x^T (I-M)^k x \right) + \frac{1}{2}\sum_{k\geq 1} \mathrm{Tr}((I-M)^k)/k \right) \middle| \mathcal{E} \right] \\
&= \mathbf{E}_S \left[\exp\left(-\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) - \frac{1}{2}\sum_{k\geq 2} \left(x^T(I-M)^k x - \mathrm{Tr}((I-M)^k)/k \right) \right) \middle| \mathcal{E} \right] \\
&= \mathbf{E}_S \left[\exp\left(-\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) - R(x) \right) \middle| \mathcal{E} \right],
\end{aligned} \tag{20}$$

where $R(x) := \frac{1}{2}\sum_{k\geq 2} \left(x^T(I-M)^k x - \mathrm{Tr}((I-M)^k)/k \right)$.

The rationale behind the definition of $\mathcal{E}(\gamma)$ is that for all $S \in \mathcal{E}(\gamma)$ the residual $R(x)$ above is (essentially) dominated by the quadratic terms, i.e. $||I-M||_F^2$ and $x^T(I-M)^2x$ (for 'typical' values of x – see Lemma 20 below), i.e. we can truncate the Taylor expansion to the first and second terms and control the error. This is made formal by the following three lemmas.

Lemma 18. *For every $\gamma \in (0,1)$, conditioned on $\mathcal{E}(\gamma)$ we have $\mathrm{Tr}((I-M)^k) \leq \gamma^{k-2} \cdot ||I-M||_F^2$ for all $k \geq 2$.*

Proof. $|\mathrm{Tr}((I-M)^k)| \leq ||I-M||_2^{k-2} \cdot \mathrm{Tr}((I-M)^2) \leq ||I-M||_F^{k-2} \cdot ||I-M||_F^2 \leq \gamma^k$ as required, since $||A||_2 \leq ||A||_F$ and $\mathrm{Tr}(A^TA) = ||A||_F^2$ for all $A \in \mathbb{R}^{d\times d}$. □

Lemma 19. *For any matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, any $\gamma \in (0,1/2)$, for any $x \in \mathbb{R}^d$ one has, for any CountSketch matrix $S \in \mathcal{E}(\gamma)$, $x^T(I-M)^k x \leq \gamma^{k-2} x^T(I-M)^2 x$ for any $k \geq 2$.*

Proof. We have, for any $x \in \mathbb{R}^d$ and any $S \in \mathcal{E}(\gamma)$ $|x^T(I-M)^k x| \leq ||I-M||_2^{k-2} \cdot x^T(I-M)^2 x \leq \gamma^{k-2} \cdot x^T(I-M)^2 x$, as $||I-M||_2 \leq ||I-M||_F$.

□

Lemma 20. *For any $\gamma \in (0,1/2)$, any matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, any CountSketch matrix $S \in \mathcal{E}(\gamma)$ and any $x \in \mathcal{T}(S,U)$ one has*

$$|R(x)| \leq \sum_{k\geq 2} |x^T(I-M)^k x| + |\mathrm{Tr}((I-M)^k)|/k \leq C||I-M||_F^2 + Cx^T(I-M)^2x,$$

where $C > 0$ is an absolute constant.

Proof. We have by combining Lemma 18 and Lemma 19

$$\begin{aligned}
\sum_{k\geq 2} |x^T(I-M)^k x| + |\mathrm{Tr}((I-M)^k)|/k &\leq \sum_{k\geq 2} [\gamma^{k-2} x^T(I-M)^2 x + \gamma^{k-2} \cdot ||I-M||_F^2/k] \\
&\leq C(x^T(I-M)^2x + ||I-M||_F^2)
\end{aligned}$$

for an absolute constant $C' > 0$, as $\gamma < 1/2$ by assumption of the lemma. □

7.4 Constructing the universal set $\mathcal{T}^*(U,\gamma)$ of typical x

The main result of this section is the following lemma:

Lemma 21. *For every matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, for every $\gamma \in (0, 1/\log^2 n)$ and any $\delta > 0$ if*

$$\begin{aligned}\mathcal{T}^*(U,\gamma) := \big\{x \in \mathbb{R}^d \text{ s.t. } &||x||_\infty \le C\sqrt{\log n} \text{ and}\\ &|(Ux)_a| \le O(\sqrt{\log n})||U_a||_2 \text{ for all } a \in [n] \text{ and}\\ &\mathbf{E}_S\left[\mathbf{I}[x \notin \mathcal{T}(S,U)]|\,\mathcal{E}(\gamma)\right] < 1/n^{25}.\big\},\end{aligned}$$

then **(a)** $\mathbf{Pr}_{X\sim N(0,I_d)}[X \in \mathcal{T}^*(U,\gamma)] \ge 1 - n^{-10}$ *and* **(b)** $\mathbf{Pr}_{X\sim\tilde{q}}[X \in \mathcal{T}^*(U,\gamma)] \ge 1 - n^{-10}$.

Note that the lemma guarantees the existence of a universal set $\mathcal{T}^* \subseteq \mathbb{R}^d$ that captures most of the probability mass of both the normal distribution $N(0, I_d)$ and the mixture $\tilde{q}$.

Proof of Lemma 21:

Let

$$\mathcal{T}_1^* := \{x \in \mathbb{R}^d : \mathbf{E}_S\left[\mathbf{I}[x \notin \mathcal{T}(S,U)]|\,\mathcal{E}(\gamma)\right] < 1/n^{25}\}.$$

$$\mathcal{T}_2^* := \{x \in \mathbb{R}^d : ||x||_\infty \le C\sqrt{\log n}\}.$$

$$\mathcal{T}_3^* := \{x \in \mathbb{R}^d : |(Ux)_a| \le C\sqrt{\log n}||U_a||_2 \text{ for all } a \in [n]\}.$$

We prove that $\mathcal{T}_i^*, i = 1, 2, 3$ occur with high probability under both distributions. As we show below, the result then follows by a union bound.

Showing that $\mathcal{T}_1^*$ occurs with high probability. We first show that $\mathcal{T}_1^*$ occurs with high probability under the isotropic Gaussian distribution $X \sim N(0, I_d)$, and then show that it also occurs with high probability under the mixture of Gaussians distribution $\tilde{q}$. In both cases the proof proceeds by applying Claim 16 followed by Markov's inequality.

Step 1: bounding $\mathbf{Pr}_{X\sim N(0,I_d)}[\mathcal{T}_1^*]$. We have by Claim 16, **(1)** that $\mathbf{Pr}_{X\sim N(0,I_d)}[\mathbf{I}[X \notin \mathcal{T}(S,U)]] < n^{-40}$, and hence

$$\mathbf{E}_S\left[\mathbf{E}_{X\sim N(0,I_d)}\left[\mathbf{I}[X \notin \mathcal{T}(S,U)]\right]\big|\,\mathcal{E}(\gamma)\right] < 1/n^{40},$$

implying that $\mathbf{E}_{X\sim N(0,I_d)}\left[\mathbf{E}_S\left[\mathbf{I}[X \notin \mathcal{T}(S,U)]\right]|\,\mathcal{E}(\gamma)\right] < 1/n^{40}$. We thus get by Markov's inequality that $\mathbf{Pr}_{X\sim N(0,I_d)}[\mathcal{T}_1^*] \ge 1 - n^{-15}$.

Step 2: bounding $\mathbf{Pr}_{X\sim\tilde{q}}[\mathcal{T}_1^*]$. We have by Claim 16, **(2)** that for any $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, any pair of matrices $S, S' \in \mathcal{E}(\gamma)$, if $M' = U^T S^T S U$, then $\mathbf{Pr}_{X\sim N(0,M')}[X \notin \mathcal{T}(S,U)] < n^{-40}$. We thus have

$$\begin{aligned}\mathbf{E}_{X\sim\tilde{q}}[\mathbf{E}_S[\mathbf{I}[X \notin \mathcal{T}(S,U)]|\mathcal{E}(\gamma)]] &= \mathbf{E}_{S'}\left[\mathbf{E}_{X\sim q_{S'}}[\mathbf{E}_S[\mathbf{I}[X \notin \mathcal{T}(S,U)]]|\mathcal{E}(\gamma)]\big|\,\mathcal{E}(\gamma)\right]\\ &= \mathbf{E}_S\left[\mathbf{E}_{S'}[\mathbf{E}_{X\sim q_{S'}}[\mathbf{I}[X \notin \mathcal{T}(S,U)]]|\mathcal{E}(\gamma)]\big|\,\mathcal{E}(\gamma)\right]\\ &= \mathbf{E}_S\left[\mathbf{Pr}_{X\sim\tilde{q}}[\mathbf{I}[X \notin \mathcal{T}(S,U)]]|\,\mathcal{E}(\gamma)\right]\\ &\le n^{-40}.\end{aligned}$$

By Markov's inequality applied to the expression in the first line we thus have

$$\mathbf{Pr}_{X\sim\tilde{q}}[\mathbf{E}_S[\mathbf{I}[X \notin \mathcal{T}(S,U)]|\mathcal{E}(\gamma)] > n^{-25}] < n^{-15}.$$

Showing that $\mathcal{T}_2^*$ occurs with high probability. The fact that

$$\mathbf{Pr}_{X\sim N(0,I_d)}\left[||X||_\infty \le C\sqrt{\log n}\right] \ge 1 - n^{-40}$$

follows by standard properties of Gaussian random variables. Thus, it remains to show that $\mathcal{T}_2^*$ occurs with high probability under $X \sim \tilde{q}$. For any $U \in \mathbb{R}^{n\times d}$ and $S \in \mathcal{E}(\gamma)$ we now prove that for $M = U^T S^T S U$

$$\mathbf{Pr}_{X\sim N(0,M)}\left[||X||_\infty \le C\sqrt{\log n}\right] \ge 1 - n^{-40} \tag{21}$$

It is convenient to let $X = M^{1/2}Y$, where $Y \sim N(0, I_d)$ is a vector of independent Gaussians of unit variance. Then we need to bound

$$\mathbf{Pr}_{X\sim N(0,M)}\left[||X||_\infty \ge C\sqrt{\log n}\right] = \mathbf{Pr}_{Y\sim N(0,I_d)}\left[||M^{1/2}Y||_\infty \ge C\sqrt{\log n}\right]$$

By 2-stability of the Gaussian distribution we have that for each $i = 1, \ldots, d$ the random variable $(M^{1/2}Y)_i$ is Gaussian with variance at most $||M^{1/2}||_F^2$, which we bound by

$$\begin{aligned}
||M^{1/2}||_F = ||(I + (M-I))^{1/2}||_F &= \left|\left|\sum_{t=0}^{\infty}\binom{1/2}{t}(I-M)^t\right|\right|_F \\
&\le \sum_{t=0}^{\infty}\left|\binom{1/2}{t}\right| \cdot \left|\left|(I-M)^t\right|\right|_F \\
&\le \sum_{t=0}^{\infty}\left|\binom{1/2}{t}\right| \cdot ||I-M||_F^t \\
&\le \sum_{t=0}^{\infty}||I-M||_F^t \\
&\le \sum_{t=0}^{\infty}(1/2)^t \\
&\le 2
\end{aligned}$$

Thus, for each $i \in [n]$ the random variable $(M^{1/2}Y)_i$ is Gaussian with variance at most 4, and (21) follows by standard properties of Gaussian random variables as long as $C > 0$ is a sufficiently large constant.

Showing that $\mathcal{T}_3^*$ occurs with high probability. The fact that

$$\mathbf{Pr}_{X\sim N(0,I_d)}\left[|(UX)_a| \le C\sqrt{\log n}\cdot||U_a||_2 \text{ for all } a \in [n]\right] \ge 1 - n^{-40}$$

follows by standard properties of Gaussian random variables and a union bound over all $a \in [n]$.

Thus, it remains to show that $\mathcal{T}_3^*$ occurs with high probability under $X \sim \tilde{q}$. For any $U \in \mathbb{R}^{n\times d}$ and $S \in \mathcal{E}(\gamma)$ we now prove that for $M = U^T S^T S U$

$$\mathbf{Pr}_{X\sim N(0,M)}\left[|(UX)_a| \le C\sqrt{\log n}||U_a||_2 \text{ for all } a \in [n]\right] \ge 1 - n^{-40}$$

It is convenient to let $X = M^{1/2}Y$, where $Y \sim N(0, I_d)$ is a vector of independent Gaussians of unit variance. Then we need to bound, for each $a \in [n]$

$$\mathbf{Pr}_{X\sim N(0,M)}\left[|(UX)_a| \ge C\sqrt{\log n}||U_a||_2\right] = \mathbf{Pr}_{Y\sim N(0,I_d)}\left[|(UM^{1/2}Y)_a| \ge C\sqrt{\log n}||U_a||_2\right]$$

By 2-stability of the Gaussian distribution we have that for each $a = 1, \ldots, n$ the random variable $U_a M^{1/2}Y$ is Gaussian with variance at most $||U_a M^{1/2}||_2^2 \le 4||U_a||_F^2$ (since $\gamma < 1/\log^2 n$ by assumption of the lemma), and hence the result follows by standard properties of Gaussian random variables and a union bound.

Finally, we let $\mathcal{T}^* := \mathcal{T}_1^* \cap \mathcal{T}_2^* \cap \mathcal{T}_3^*$. By a union bound applied to the bounds above we have that $\mathcal{T}^*$ occurs with probability at least $1 - n^{-10}$ under both distributions, as required. $\square$

7.5 Proof of Lemma 10

We first prove

Lemma 22. *There exists an absolute constant $C > 0$ such that for every $\gamma \in (0, 1/\log n)$, any matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns and any CountSketch matrix $S \in \mathcal{E}(\gamma)$ and $x \in \mathcal{T}(S, U)$ one has, letting*

$$L(x) := -\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) - \frac{1}{8}x^T(I-M)x \cdot \mathrm{Tr}(I-M)$$

and

$$Q(x) := ((x^T(I-M)x)^2 + (\mathrm{Tr}(I-M))^2 + x^T(I-M)^2x + ||I-M||_F^2),$$

that

$$\left|1 + L(x) - \exp\left(\frac{1}{2}x^Tx - \frac{1}{2}x^TM^{-1}x - \frac{1}{2}\log\det M\right)\right| \le C \cdot Q(x).$$

The proof is given in section 8.

We will need the following two lemmas, whose proofs are provided in section 8.2

Lemma 23. *For any $U \in \mathbb{R}^{n\times d}$ with orthonormal columns, and $B \ge 1$, if S is a random CountSketch matrix and $M = U^TS^TSU$, then*

(1) $\mathbf{E}_S[||M-I||_F^2] \le 2d^2/B$

(2) *for all $x \in \mathcal{T}^*$ one has* $\mathbf{E}_S[x^T(I-M)^2x] = O(d^2(\log^2 n)/B)$

(3) *for all $x \in \mathcal{T}^*$ one has* $\mathbf{E}_S[(x^T(I-M)x)^2] = O(d^2(\log^2 n)/B)$

(4) *for all $x \in \mathcal{T}^*$ one has* $\mathbf{E}_S[(x^T(I-M)x)\cdot\mathrm{Tr}(I-M)] = O(d^2(\log n)/B)$

(5) *one has* $\mathbf{E}_S[(\mathrm{Tr}(I-M))^2] = O(d^2/B)$

and

Lemma 24 (Variance bound)**.** *For any matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns if $\gamma \in (0, 1/2)$ and $\mathcal{T}^*(U,\gamma) \subseteq \mathbb{R}^d$ is as defined in Lemma 21, then for any $x \in \mathcal{T}^*(U,\gamma)$ one has, for*

$$L(x) := -\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) - \frac{1}{8}x^T(I-M)x \cdot \mathrm{Tr}(I-M)$$

and

$$Q(x) := ((x^T(I-M)x)^2 + (\mathrm{Tr}(I-M))^2 + x^T(I-M)^2x + ||I-M||_F^2),$$

that for any constant C

$$\mathbf{E}_S\left[(L(x) + C\cdot Q(x))^2\right] = O(d^2(\log^2 n)/B),$$

where S is a uniformly random CountSketch matrix and $M = U^TS^TSU$.

We will use the following lemma, whose proof is given in section 8:

Lemma 25. *For any random variable Z and any event $\mathcal{E}$ with $\Pr[\mathcal{E}] \ge 1/2$, if $\epsilon := \mathbf{E}[(Z-1)^2]$, then*

$$|\mathbf{E}[Z] - \mathbf{E}[Z|\mathcal{E}]| \le 2(1+\mathbf{E}[Z])\mathbf{Pr}[\bar{\mathcal{E}}] + 2\sqrt{\epsilon\mathbf{Pr}[\bar{\mathcal{E}}]}.$$

Equipped with the bounds above, we can now prove Lemma 10:

Lemma 10 (Restated) *There exists an absolute constant $C > 0$ such that for every $\delta \in (0,1)$ and every matrix $U \in \mathbb{R}^{n\times d}$ with orthonormal columns if $B \geq \frac{1}{\delta}C(\log n)^4 d^2$ there exists a set $\mathcal{E}$ of CountSketch matrices and a subset $\mathcal{T}^* \subseteq \mathbb{R}^d$ that satisfies $\mathbf{Pr}_{X\sim p}[X \notin \mathcal{T}^*] \leq n^{-10}$ and $\mathbf{Pr}_{X\sim \tilde{q}}[X \notin \mathcal{T}^*] \leq n^{-10}$ such that if $S \in \mathbb{R}^{B\times n}$ is a random CountSketch matrix, then* **(1)** $\mathbf{Pr}_S[\mathcal{E}] \geq 1 - \delta/3$*, and* **(2)** *for all $x \in \mathcal{T}^*$ one has*

$$\left|\frac{q'(x)}{p(x)} - 1\right| \leq O((d^2 \log^4 n)/B) + O(n^{-10}).$$

Proof. Let $\mathcal{T}^*(U,\gamma) \subseteq \mathbb{R}^d$ be as defined in Lemma 21, and let $\gamma := 1/\log^2 n$. Let $\mathcal{E} := \mathcal{E}(\gamma)$, and note that $\mathbf{Pr}[\mathcal{E}] \geq 1 - \delta/3$ by Claim 13 as long as C is a large enough constant, as required.

We now bound

$$\frac{q'(x)}{p(x)} = \mathbf{E}_S\left[\frac{q_S(x)}{p(x)} \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] + \xi,$$

for $x \in \mathcal{T}^*(U,\gamma)$, where $\xi = \mathbf{E}_S[\mathbf{Pr}_{X\sim q_S}[X \in \mathcal{T}(S,U)]] \leq n^{-40}$ by definition and Claim 17, **(2)**. For each $S \in \mathcal{E}(\gamma)$ and $x \in \mathcal{T}(S,U)$ we have by Lemma 22

$$\left|\frac{q_S(x)}{p(x)} - (1 + L(x))\right| = \left|\exp\left(\frac{1}{2}x^T x - \frac{1}{2}x^T M^{-1}x - \frac{1}{2}\log\det M\right) - (1 + L(x))\right| \leq C \cdot Q(x),$$

where

$$L(x) := -\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) - \frac{1}{8}x^T(I-M)x \cdot \mathrm{Tr}(I-M)$$

denotes the 'linear' term and

$$Q(x) := (x^T(I-M)x)^2 + (\mathrm{Tr}(I-M))^2 + x^T(I-M)^2x + ||I-M||_F^2$$

denotes the 'quadratic' term.

Taking expectations, we get

$$\begin{aligned}
&\mathbf{E}_S\left[(L(x) - C\cdot Q(x)) \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] \\
&\leq \mathbf{E}_S\left[\left(\exp\left(\frac{1}{2}x^T x - \frac{1}{2}x^T M^{-1}x - \frac{1}{2}\log\det M\right) - 1\right) \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] \\
&\leq \mathbf{E}_S\left[(L(x) + C\cdot Q(x)) \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right].
\end{aligned}$$

Thus, it suffices to show that

$$|\mathbf{E}_S\left[(L(x) \pm C\cdot Q(x)) \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right]| = O((Cd\log n)^2/B) + O(n^{-10}),$$

which we do now. We only provide the analysis for the case when the sign in front of the constant C is a plus, as the other part is analogous.

We first show that removing the multiplier $\mathbf{I}[x \in \mathcal{T}(S,U)]$ from the equation above only changes the expectation slightly. Specifically, note that

$$\begin{aligned}
&|\mathbf{E}_S\left[(L(x) + C\cdot Q(x)) \cdot \mathbf{I}[x \in \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] - \mathbf{E}_S\left[L(x) + C\cdot Q(x) \middle| \mathcal{E}(\gamma)\right]| \\
&\leq \mathbf{E}_S\left[|L(x) + C\cdot Q(x)| \cdot \mathbf{I}[x \notin \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right].
\end{aligned} \tag{22}$$

By Claim 15 we have $||I-M||_F^2 \leq 4n^3$ for all S and U, so every element of the matrix $I-M$ is upper bounded by $2n^2$. Similarly, we have $||(I-M)^2||_F \leq ||I-M||_F^2$, and so every element of $(I-M)^2$ is upper bounded by $4n^3$. Thus, for any $x \in \mathcal{T}^*(U,\gamma)$ one has

$$\begin{aligned}
&|L(x) + CQ(x)| \\
&\leq (|x^T(I-M)x| + |\mathrm{Tr}(I-M)| + |x^T(I-M)x \cdot \mathrm{Tr}(I-M)| \\
&+ C((x^T(I-M)x)^2 + (\mathrm{Tr}(I-M))^2 + x^T(I-M)^2x + ||I-M||_F^2)) \\
&= O(\log n)(2n^2d^2 + d\cdot(2n^2) + (2n^2)^2d^3 + (2n^2d^2)^2 + (d\cdot 2n^2)^2 + 4n^4d^2 + 4n^3) \leq n^{10}
\end{aligned}$$

as long as n is sufficiently large, where we used the fact that $||x||_\infty \le O(\sqrt{\log n})$ for all $x \in \mathcal{T}^*(U,\gamma)$.

Furthermore, by Lemma 21 we have for $x \in \mathcal{T}^*(U,\gamma)$ that

$$\mathbf{E}_S\left[\mathbf{I}[x \notin \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] < 1/n^{25}.$$

Substituting these two bounds into (22), we get

$$\mathbf{E}_S\left[|L(x) + C\cdot Q(x)| \cdot \mathbf{I}[x \notin \mathcal{T}(S,U)] \middle| \mathcal{E}(\gamma)\right] \le n^{-10} \tag{23}$$

so it remains to bound

$$\mathbf{E}_S\left[L(x) + C\cdot Q(x) \middle| \mathcal{E}(\gamma)\right].$$

We bound the expectation above by relating it to the corresponding unconditional expectation. Let $Z := 1 + (L(x) + C\cdot Q(x))$, and note that

$$\mathbf{E}_S[Z] = 1 - \mathbf{E}_S[\frac{1}{8}x^T(I-M)x \cdot \mathrm{Tr}(I-M)] + C\cdot \mathbf{E}_S[Q(x)] = 1 + O((C\log n)^2 d^2/B) \tag{24}$$

by Lemma 23. Let $\epsilon := \mathbf{E}_S[(Z-1)^2]$. We note that by Lemma 24 that $\epsilon \le O(d^2(\log^2 n)/B)$, and hence since $\mathcal{E}(\gamma) \ge 1/2$ by Claim 13, by Lemma 25 we have

$$|\mathbf{E}[Z] - \mathbf{E}[Z|\mathcal{E}(\gamma)]| \le 2(1+\mathbf{E}[Z])\mathbf{Pr}[\bar{\mathcal{E}}(\gamma)] + 2\sqrt{\epsilon \mathbf{Pr}[\bar{\mathcal{E}}(\gamma)]}.$$

Since $\mathbf{Pr}[\bar{\mathcal{E}}(\gamma)] \le 3(d/\gamma)^2/B$ by Claim 13 and using the assumption that $B \ge (\log^2 n)d^2$, we get

$$|\mathbf{E}[Z] - \mathbf{E}[Z|\mathcal{E}(\gamma)]| \le O((d/\gamma)^2/B) + 2\sqrt{O(d^2\log^2 n/B)\cdot(d/\gamma)^2/B)} = O((\frac{1}{\gamma^2} + \frac{1}{\gamma}\log n)d^2/B) = O((d/\gamma)^2/B), \tag{25}$$

where we used the assumption that $\gamma \le 1/\log^2 n$. Combining (25), (24) with (22) and (23), we get

$$\left|\frac{q'(x)}{p(x)} - 1\right| = \left|\mathbf{E}_S\left[\frac{q(x)}{p(x)}\cdot \mathbf{I}[x \in \mathcal{T}(U,S)] \middle| \mathcal{E}(\gamma)\right] + \xi - 1\right| \le O((d^2\log^4 n)/B) + O(1/n^{10}).$$

□

8 Proofs omitted from the main body

8.1 Proof of Claim 16 and Claim 15

We will use

Theorem 26 (Bernstein's inequality)**.** *Let $X_1, \ldots, X_n$ be independent zero mean random variables such that $|X_i| \le L$ for all i with probability 1, and let $X := \sum_{i=1}^n X_i$. Then*

$$\mathbf{Pr}[X > t] < \exp\left(-\frac{\frac{1}{2}t^2}{\sum_{i=1}^n \mathbf{E}[X_i^2] + \frac{1}{3}Lt}\right).$$

Proof of Claim 16:

Proving (1). The bound follows by standard concentration inequalities, as we now show. Since the normal distribution is rotationally invariant, we have that

$$X^T(I-M)X = \sum_{i=1}^{d}(\lambda_i - 1)Y_i^2 = \mathrm{Tr}(M-I) + \sum_{i=1}^{d}(\lambda_i-1)(Y_i^2-1), \tag{26}$$

where $Y \sim N(0, I_d)$ and λ_i are the eigenvalues of M. We now apply Bernstein's inequality (Theorem 26) to random variables $(\lambda_i - 1)(Y_i^2 - 1)$ (note that they are zero mean). We also have $\mathbf{E}[(\lambda_i-1)^2(Y_i^2-1)^2] \leq O((\lambda_i-1)^2)$. We later combine it with the fact that $|\mathrm{Tr}(I-M)| \leq \gamma \leq \frac{1}{2}\cdot\frac{1}{100}$ for all $S \in \mathcal{E}(\gamma)$ to obtain the result. We also have $|(\lambda_i-1)Y_i| \leq ||I-M||_F C\sqrt{\log n} \leq \gamma \cdot C\sqrt{\log n}$ for all i with probability at least $1 - n^{-40}/4$ as long as $C > 0$ is larger than an absolute constant. We thus have by applying Theorem 26 to random variables clipped at $\gamma C\sqrt{\log n}$ in magnitude, which we denote by event $\mathcal{F}$, to conclude for all $t \geq 0$,

$$\mathbf{Pr}[\sum_{i=1}^{d}(\lambda_i-1)(Y_i^2-1) > t \mid \mathcal{F}] < 2\exp\left(-\frac{\frac{1}{2}t^2}{O(\sum_{i=1}^{n}(\lambda_i-1)^2) + (\frac{1}{3}\gamma C\sqrt{\log n})t}\right).$$

Note the random variables are still independent and zero-mean conditioned on $\mathcal{F}$, and $\mathbf{E}[(\lambda_i-1)^2(Y_i^2-1)^2] \leq O((\lambda_i-1)^2)$ continues to hold, since the clipping changes the expectation by at most a factor of $(1+O(n^{-40}))$. By a union bound we can remove the conditioning on $\mathcal{F}$,

$$\mathbf{Pr}[\sum_{i=1}^{d}(\lambda_i-1)(Y_i^2-1) > t] < 2\exp\left(-\frac{\frac{1}{2}t^2}{O(\sum_{i=1}^{n}(\lambda_i-1)^2) + (\frac{1}{3}\gamma C\sqrt{\log n})t}\right) + \frac{n^{-40}}{4}.$$

Setting $t = \frac{1}{100}$, and using the fact that $\sum_i(\lambda_i-1)^2 = ||I-M||_F^2 \leq \gamma^2$, we get

$$\begin{aligned}
\mathbf{Pr}[\sum_{i=1}^{d}(\lambda_i-1)(Y_i^2-1) > \frac{1}{2}\cdot\frac{1}{100}] &< 2\exp\left(-\frac{\frac{1}{2}(\frac{1}{2}\cdot\frac{1}{100})^2}{O(\gamma^2) + (\frac{1}{3}\cdot(\frac{1}{2}\cdot\frac{1}{100})\gamma C\sqrt{\log n})}\right) + \frac{n^{-40}}{4} \\
&= \exp(-\Omega(1/(\gamma\sqrt{\log n}))) + \frac{n^{-40}}{4} \\
&< \frac{n^{-40}}{2},
\end{aligned}$$

since $\gamma \leq 1/\log^2 n$ by assumption, for a sufficiently large n. Combining this with (26), we get, using the fact that $|\mathrm{Tr}(I-M)| \leq \gamma < \frac{1}{2}\cdot\frac{1}{100}$ for $S \in \mathcal{E}(\gamma)$ that

$$\mathbf{Pr}[X^T(I-M)X > \frac{1}{100}] \leq \mathbf{Pr}[|\sum_{i=1}^{d}(\lambda_i-1)(Y_i^2-1)| > \frac{1}{2}\frac{1}{100}] < n^{-40}/2,$$

as required.

We also have

$$X^T(I-M)^2X = \sum_{i=1}^{d}(\lambda_i-1)^2Y_i^2 \leq ||I-M||_F^2 \cdot \max_{i\in[d]}|Y_i|^2 \leq O(\log n)\cdot||I-M||_F^2 = O(\log n\gamma^2) \leq \frac{1}{100}$$

with probability at least $1 - n^{-40}/2$ by standard properties of Gaussian random variables. Putting the two estimates together and taking a union bound over the failure events now shows that $\mathbf{Pr}_{X\sim N(0,I_d)}[X \notin \mathcal{T}(S,U)] < n^{-40}$, as required.

Proving (2). Recall that $\mathcal{T}(S,U) = \left\{x \in \mathbb{R}^d : |x^T(I-M)x| \leq \frac{1}{100} \text{ and } x^T(I-M)^2x \leq \frac{1}{100}\right\}$. For any S' we have that $X \sim N(0,M')$, where $M' = (S'U)^TS'U$, so $X = M'^{1/2}Y$, where $Y = N(0,I_d)$. We thus have

$$X^T(I-M)X = (M'^{1/2}Y)^T(I-M)(M'^{1/2}Y) = Y^TM'^{1/2}(I-M)M'^{1/2}Y.$$

We now show that

$$\mathbf{Pr}_{Y\sim N(0,I_d)}\left[\left|Y^TM'^{1/2}(I-M)M'^{1/2}Y\right| > \frac{1}{100}\right] < 1/n^{20} \tag{27}$$

Let $Q := M'^{1/2}(I-M)M'^{1/2}$, and let $1-\tilde{\lambda}_i, i = 1,\ldots,d$ denote the eigenvalues of Q. We have

$$Y^TM'^{1/2}(I-M)M'^{1/2}Y = \sum_{i=1}^d(1-\tilde{\lambda}_i)Z_i^2,$$

where $Z \sim N(0,I_d)$. Note that

$$\begin{aligned}
|\sum_{i=1}^d(1-\tilde{\lambda}_i)| &= |\mathrm{Tr}(Q)| = |\mathrm{Tr}(M'^{1/2}(I-M)M'^{1/2})| \\
&= |\mathrm{Tr}(M'(I-M))| = |\mathrm{Tr}((I-(I-M'))(I-M))| \\
&\leq |\mathrm{Tr}(I-M)| + |\mathrm{Tr}\left((I-M')\left(I-M\right)\right)| \\
&= \gamma + |\mathrm{Tr}\left((I-M')\left(I-M\right)\right)| \quad \text{(since } |\mathrm{Tr}(I-M)| \leq \gamma \text{ for all } S \in \mathcal{E}(\gamma)) \\
&\leq \gamma + ||I-M'||_F \cdot ||M-I||_F \quad \text{(by von Neumann and Cauchy-Schwarz inequalities)} \\
&\leq \gamma + \gamma^2
\end{aligned} \tag{28}$$

We thus have

$$\begin{aligned}
Y^TM'^{1/2}(I-M)M'^{1/2}Y &= \sum_{i=1}^d(1-\tilde{\lambda}_i)Z_i^2 \\
&= \sum_{i=1}^d(1-\tilde{\lambda}_i) + \sum_{i=1}^d(1-\tilde{\lambda}_i)(Z_i^2-1)
\end{aligned} \tag{29}$$

We now use a calculation analogous to the above for **(1)** to show that $|\sum_{i=1}^d(1-\tilde{\lambda}_i)(Z_i^2-1)| \leq \frac{1}{2}\cdot\frac{1}{100}$ with probability at least $1-n^{-40}/4$. Indeed, we first verify that the variance is bounded by

$$\begin{aligned}
O(\sum_{i=1}^d(1-\tilde{\lambda}_i)^2) &= O(||Q||_F^2) \\
&= O(||M'^{1/2}(I-M)M'^{1/2}||_F^2) \\
&\leq O(||M'||_2^2)||I-M||_F^2 \quad \text{(by sub-multiplicativity)} \\
&\leq O((||I||_2 + ||M'-I||_F)^2)||I-M||_F^2 \\
&\leq O(||I-M||_F^2) \\
&= O(\gamma^2).
\end{aligned} \tag{30}$$

We also have

$$\begin{aligned}
|(1-\tilde{\lambda}_i)Y_i| &\leq ||Q||_FC\sqrt{\log n} \\
&\leq ||M'||_2||I-M||_FC\sqrt{\log n} \quad \text{(by sub-multiplicativity)} \\
&\leq (||I||_2 + ||M'-I||_F)||I-M||_FC\sqrt{\log n} \\
&\leq 2||I-M||_FC\sqrt{\log n} \\
&\leq 2\gamma\cdot C\sqrt{\log n},
\end{aligned}$$

for all i with probability at least $1 - 1/n^{40}/5$ as long as $C > 0$ is larger than an absolute constant. We thus have by Theorem 26 (applied to clipped variables and then unclipping by a union bound as in (**1**)) for all $t \geq 0$ that

$$\mathbf{Pr}[|Y^T M'^{1/2}(I-M)M'^{1/2}Y - \sum_{i=1}^{d}(1-\tilde{\lambda}_i)| > t] < \exp\left(-\frac{\frac{1}{2}t^2}{O(\sum_{i=1}^{n}(1-\tilde{\lambda}_i)^2) + (\frac{1}{3}2\gamma C\sqrt{\log n})t}\right) + n^{-40}/5.$$

Setting $t = \frac{1}{2}\frac{1}{100}$, and using the upper bound $O(\sum_i (1-\tilde{\lambda}_i)^2) = O(\gamma^2)$ obtained in (30), we get

$$\begin{aligned}\mathbf{Pr}[|Y^T M'^{1/2}(I-M)M'^{1/2}Y - \sum_{i=1}^{d}(1-\tilde{\lambda}_i)| > \frac{1}{2}\cdot\frac{1}{100}] &< \exp\left(-\frac{\frac{1}{2}(\frac{1}{2}\cdot\frac{1}{100})^2}{C\gamma^2 + (\frac{1}{3}\cdot\frac{1}{2}\frac{1}{100}\gamma C\sqrt{\log n})}\right) + n^{-40}/5 \\ &= \exp(-\Omega(1/(\gamma\sqrt{\log n}))) + n^{-40}/5 < n^{-40}/4\end{aligned}$$

since $\gamma \leq 1/\log^2 n$ by assumption, for a sufficiently large n. Since $|\sum_{i=1}^{d}(1-\tilde{\lambda}_i)| \leq \gamma + 2\gamma^2 \leq \frac{1}{2}\cdot\frac{1}{100}$ by (28), we get by triangle inequality that

$$\mathbf{Pr}_{X\sim N(0,M')}[|X^T(I-M)X| > \frac{1}{100}] \leq n^{-40}/4.$$

Similarly to **(1)** above, we have, when $X \sim N(0, M')$, $X = M'^{1/2}Y, Y \sim N(0, I_d)$,

$$\begin{aligned}X^T(I-M)^2X = Y^T M'^{1/2}(I-M)^2M'^{1/2}Y &= \sum_{i=1}^{d}\tilde{\tau}_i Z_i^2 \\ &\leq \mathrm{Tr}(M'^{1/2}(I-M)^2M^{1/2})\cdot\max_{i\in[d]}Z_i^2 \\ &\leq O(\log n)\cdot\mathrm{Tr}(M'^{1/2}(I-M)^2M'^{1/2})\end{aligned}$$

with probability at least $1 - n^{-40}/2$ over the choice of X, as $\max_{i\in[d]} Z_i^2 \leq C\log n$ with high probability if C is a sufficiently large constant by standard properties of Gaussian random variables. Since $\mathrm{Tr}(M'^{1/2}(I-M)^2M'^{1/2}) = \mathrm{Tr}(M'(I-M)^2) \leq 2||I-M||_F^2$ (as $\gamma < 1/\log^2 n < 1/3$ by assumption of the lemma), we get

$$X^T(I-M)^2X \leq O(\log n)\cdot\mathrm{Tr}(M'^{1/2}(I-M)^2M'^{1/2}) \leq O(\log n)\cdot\gamma^2 \leq \frac{1}{100} \quad (\text{since } \gamma < 1/\log^2 n)$$

with probability at least $1 - n^{-40}/4$. A union bound over the failure events yields $\mathbf{Pr}_{X\sim N(0,M')}[X \notin \mathcal{T}(S,U)] < n^{-40}$, as required.

This completes the proof. □

Proof of Lemma 22: By assumption that $S \in \mathcal{E}(\gamma)$ we have that $||I-M||_2 \leq \gamma$, so Taylor expansion is valid and gives

$$\frac{1}{2}x^Tx - \frac{1}{2}x^TM^{-1}x - \frac{1}{2}\log\det M = -\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) + R(x),$$

where for all $x \in \mathcal{T}(S,U)$ one has $R(x) \leq \sum_{k\geq 2} x^T(I-M)^2x + \mathrm{Tr}(I-M)^k$.

We have by Lemma 20 that $R(x) \leq C(x^T(I-M)^2x + ||I-M||_F^2)$ for an absolute constant $C > 0$, for all $x \in \mathcal{T}(S,U)$ and $S \in \mathcal{E}(\gamma)$. We thus have

$$\begin{aligned}&e^{-\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) - C(x^T(I-M)^2x + ||I-M||_F^2)} \\ &\leq e^{-\frac{1}{2}x^Tx + \frac{1}{2}\mathrm{Tr}(I-M) - \frac{1}{2}x^TM^{-1}x - \frac{1}{2}\log\det M} \\ &\leq e^{-\frac{1}{2}x^T(I-M)x + \frac{1}{2}\mathrm{Tr}(I-M) + C(x^T(I-M)^2x + ||I-M||_F^2)}\end{aligned} \tag{31}$$

for all such M and x.

We now Taylor expand $e^{-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+A(x^T(I-M)^2x+||I-M||_F^2)}$, where A is any constant (positive or negative), getting

$$\begin{aligned}&e^{-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+A(x^T(I-M)^2x+||I-M||_F^2)}\\&=\sum_{k\geq 1}\left(-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+A(x^T(I-M)^2x+||I-M||_F^2)\right)^k/k!.\end{aligned}\tag{32}$$

For $k=2$ we have

$$\begin{aligned}&\left|\left(-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+x^T(I-M)^2x+||I-M||_F^2\right)^2/2+\frac{1}{8}x^T(I-M)x\cdot\mathrm{Tr}(I-M)\right|\\&\leq C\left((x^T(I-M)x)^2+(\mathrm{Tr}(I-M))^2+x^T(I-M)^2x+||I-M||_F^2\right),\end{aligned}\tag{33}$$

where we used the fact $|x^T(I-M)x|\leq\frac{1}{100}$ for $x\in\mathcal{T}(S,U)$ and $|\mathrm{Tr}(I-M)|\leq\gamma<\frac{1}{100}$ for $S\in\mathcal{E}(\gamma)$.

For all $k\geq 3$ we use the bound

$$\begin{aligned}&|\left(-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+x^T(I-M)^2x+||I-M||_F^2\right)^k|\\&\leq\left(|x^T(I-M)x|+\frac{1}{2}|\mathrm{Tr}(I-M)|+x^T(I-M)^2x+||I-M||_F^2\right)^k\\&\leq\left(|x^T(I-M)x|+\frac{1}{2}|\mathrm{Tr}(I-M)|+x^T(I-M)^2x+||I-M||_F^2\right)^3\\&\leq C((x^T(I-M)x)^2+(\mathrm{Tr}(I-M))^2+x^T(I-M)^2x+||I-M||_F^2),\end{aligned}\tag{34}$$

where we used the bound $|x^T(I-M)x|+\frac{1}{2}|\mathrm{Tr}(I-M)|+x^T(I-M)^2x+||I-M||_F^2\leq 1$ to go from the second line to the third, and the last line follows from the observation that every term in the expansion of

$$\left(|x^T(I-M)x|+\frac{1}{2}|\mathrm{Tr}(I-M)|+x^T(I-M)^2x+||I-M||_F^2\right)^3$$

contains either at least a square of one of the first two terms or at least one of the last two.

Substituting these bounds into (32), we get

$$\begin{aligned}&e^{-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+A(x^T(I-M)^2x+||I-M||_F^2)}\\&=\sum_{k\geq 1}\left(-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+A(x^T(I-M)^2x+||I-M||_F^2)\right)^k/k!\\&\leq-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)-\frac{1}{8}x^T(I-M)x\cdot\mathrm{Tr}(I-M)\\&+C((x^T(I-M)x)^2+x^T(I-M)^2x+\mathrm{Tr}(I-M)^2+||I-M||_F^2)\quad\text{(for a constant } C>0\text{ that may depend on } A)\\&+\sum_{k\geq 3}(A+1)^k((x^T(I-M)x)^2+\mathrm{Tr}(I-M)^2+x^T(I-M)^2x+||I-M||_F^2)/k!\\&\leq-\frac{1}{2}x^T(I-M)x+\frac{1}{2}\mathrm{Tr}(I-M)+C''(x^T(I-M)^2x+(\mathrm{Tr}(I-M))^2+x^T(I-M)x^2+||I-M||_F^2)\end{aligned}$$

for an absolute constant $C''>0$. The provides the upper bound in the claimed result. The lower bound is provided by a similar calculation, which we omit. □

Proof of Lemma 25: Since $\mathbf{E}[(Z-1)^2]\leq\epsilon$ by assumption of the lemma, for any event $\mathcal{E}$ one has $\mathbf{E}[(Z-1)^2\cdot\mathbf{I}_{\bar{\mathcal{E}}}]\leq\epsilon$, where $\mathbf{I}_{\bar{\mathcal{E}}}$ is the indicator of $\bar{\mathcal{E}}$, the complement of $\mathcal{E}$. This also means that

$$\mathbf{E}[(Z-1)^2|\bar{\mathcal{E}}]\leq\epsilon/\mathbf{Pr}[\bar{\mathcal{E}}].$$

On the other hand, by Jensen's inequality

$$\mathbf{E}[|Z-1||\bar{\mathcal{E}}] \leq \left(\mathbf{E}[(Z-1)^2|\bar{\mathcal{E}}]\right)^{1/2},$$

and putting these two bounds together we get

$$\mathbf{E}[|Z-1|\cdot\mathbf{I}[\bar{\mathcal{E}}]] = \mathbf{E}[|Z-1||\bar{\mathcal{E}}]\cdot\mathbf{Pr}[\bar{\mathcal{E}}] \leq \mathbf{Pr}[\bar{\mathcal{E}}]\cdot\left(\mathbf{E}[(Z-1)^2|\bar{\mathcal{E}}]\right)^{1/2} \leq \mathbf{Pr}[\bar{\mathcal{E}}]\cdot\left(\epsilon/\mathbf{Pr}[\bar{\mathcal{E}}]\right)^{1/2} = \sqrt{\epsilon\cdot\mathbf{Pr}[\bar{\mathcal{E}}]}.$$

This means that

$$\begin{aligned}
|\mathbf{E}[Z]-\mathbf{E}[Z|\mathcal{E}]| &\leq \left|\mathbf{E}[Z]-\frac{1}{\mathbf{Pr}[\mathcal{E}]}\mathbf{E}[Z\cdot\mathbf{I}_{\mathcal{E}}]\right| \\
&\leq \left|\mathbf{E}[Z]-\frac{1}{\mathbf{Pr}[\mathcal{E}]}\mathbf{E}[Z]+\frac{1}{\mathbf{Pr}[\mathcal{E}]}\mathbf{E}[Z\cdot\mathbf{I}_{\bar{\mathcal{E}}}]\right| \\
&\leq \mathbf{E}[Z]\left(\frac{1}{1-\mathbf{Pr}[\bar{\mathcal{E}}]}-1\right)+\left|\frac{1}{\mathbf{Pr}[\mathcal{E}]}\mathbf{E}[Z\cdot\mathbf{I}_{\bar{\mathcal{E}}}]\right| \\
&\leq \mathbf{E}[Z]\cdot 2\mathbf{Pr}[\bar{\mathcal{E}}]+2\mathbf{E}[Z\cdot\mathbf{I}_{\bar{\mathcal{E}}}] \quad (\text{since } \frac{1}{1-x}-1\leq 2x \text{ for } x\in(0,1/2)) \\
&\leq \mathbf{E}[Z]\cdot 2\mathbf{Pr}[\bar{\mathcal{E}}]+2(\mathbf{Pr}[\bar{\mathcal{E}}]+\mathbf{E}[|Z-1|\cdot\mathbf{I}_{\bar{\mathcal{E}}}]) \\
&\leq 2(1+\mathbf{E}[Z])\mathbf{Pr}[\bar{\mathcal{E}}]+2\sqrt{\epsilon\mathbf{Pr}[\bar{\mathcal{E}}]}.
\end{aligned}$$

□

8.2 Proofs of moment bounds (Lemma 23 and Lemma 24)

Proof of Lemma 23 and Lemma 24: We start by noting that for every $i,j\in[1:d]$ the matrix $M=U^TS^TSU$ satisfies

$$\begin{aligned}
M_{ij} &= \sum_{r=1}^{B}\sum_{a=1}^{n}\sum_{b=1}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j} \\
&= \sum_{a=1}^{n} U_{a,i}U_{a,j}\left(\sum_{r=1}^{B}S_{r,a}^2\right)+\sum_{r=1}^{B}\sum_{a=1}^{n}\sum_{b=1,b\neq a}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j} \\
&= \delta_{i,j}+\sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j},
\end{aligned}$$

where $\delta_{i,j}$ equals 1 if $i=j$ and equals 0 otherwise. We thus have, for every $i,j\in[1:d]$, that

$$(M-I)_{ij} = \sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j},$$

which in particular means that

$$\begin{aligned}
\mathrm{Tr}(I-M) = -\sum_i (M-I)_{ii} &= -\sum_i\sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,i}, \\
&= -\sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n} S_{r,a}S_{r,b}\cdot U_aU_b^T,
\end{aligned} \tag{35}$$

(note that it immediately follows that $\mathbf{E}_S[\mathrm{Tr}(I-M)]=0$, as $\mathbf{E}_S[S_{r,a}S_{r,b}]=0$ for $a\neq b$) and

$$\begin{aligned} x^T(I-M)x &= -\sum_{ij}(M-I)_{ij}x_ix_j = -\sum_{i,j}\sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j}x_ix_j \\ &= -\sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n} S_{r,a}S_{r,b}(Ux)_a(Ux)_b \end{aligned} \tag{36}$$

(note that it immediately follows that $\mathbf{E}_S[x^T(I-M)x]=0$ for all x, as $\mathbf{E}_S[S_{r,a}S_{r,b}]=0$ for $a\neq b$).

We also have

$$(M-I)_{ij}^2 = \sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n}\sum_{r'=1}^{B}\sum_{\substack{c,d=1,\\ c\neq d}}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j}S_{r',c}U_{c,i}S_{r',d}U_{d,j}$$

and hence

$$\begin{aligned} ||I-M||_F^2 = \sum_{ij}(M-I)_{ij}^2 &= \sum_{ij}\sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n}\sum_{r'=1}^{B}\sum_{\substack{c,d=1,\\ c\neq d}}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j}S_{r',c}U_{c,i}S_{r',d}U_{d,j} \\ &= \sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n}\sum_{r'=1}^{B}\sum_{\substack{c,d=1,\\ c\neq d}}^{n} S_{r,a}S_{r,b}S_{r',c}S_{r',d}(\sum_i U_{a,i}U_{c,i})(\sum_j U_{b,j}U_{d,j}) \\ &= \sum_{r=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n}\sum_{r'=1}^{B}\sum_{\substack{c,d=1,\\ c\neq d}}^{n} S_{r,a}S_{r,b}S_{r',c}S_{r',d}\cdot U_aU_c^T\cdot U_bU_d^T \\ &= \sum_{r_1=1}^{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{r_2=1}^{B}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n} S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}\cdot U_{a_1}U_{a_2}^T\cdot U_{b_1}U_{b_2}^T \end{aligned} \tag{37}$$

We also need

$$\begin{aligned} x^T(I-M)^2x = ||(I-M)x||_2^2 &= \sum_{i=1}^{d}\left(\sum_{j=1}^{d}(I-M)_{ij}x_j\right)^2 \\ &= \sum_{i=1}^{d}\sum_{j=1}^{d}\sum_{\bar{j}=1}^{d} x_jx_{\bar{j}}\cdot\sum_{r=1}^{B}\sum_{\bar{r}=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n}\sum_{\substack{\bar{a},\bar{b}=1,\\ \bar{a}\neq\bar{b}}}^{n} S_{r,a}U_{a,i}S_{r,b}U_{b,j}\cdot S_{\bar{r},\bar{a}}U_{\bar{a},i}S_{\bar{r},\bar{b}}U_{\bar{b},\bar{j}} \\ &= \sum_{r=1}^{B}\sum_{\bar{r}=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n}\sum_{\substack{\bar{a},\bar{b}=1,\\ \bar{a}\neq\bar{b}}}^{n} S_{r,a}S_{r,b}S_{\bar{r},\bar{a}}S_{\bar{r},\bar{b}}\cdot(\sum_{i=1}^{d}U_{a,i}U_{\bar{a},i})(\sum_{j=1}^{d}U_{b,j}x_j)(\sum_{\bar{j}}U_{\bar{b},\bar{j}}x_{\bar{j}}) \\ &= \sum_{r=1}^{B}\sum_{\bar{r}=1}^{B}\sum_{\substack{a,b=1,\\ a\neq b}}^{n}\sum_{\substack{\bar{a},\bar{b}=1,\\ \bar{a}\neq\bar{b}}}^{n} S_{r,a}S_{r,b}S_{\bar{r},\bar{a}}S_{\bar{r},\bar{b}}\cdot U_aU_{\bar{a}}^T\cdot(Ux)_b(Ux)_{\bar{b}} \\ &= \sum_{r_1=1}^{B}\sum_{r_2=1}^{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n} S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}\cdot U_{a_1}U_{a_2}^T\cdot(Ux)_{b_1}(Ux)_{b_2} \end{aligned} \tag{38}$$

Bounding $\mathbf{E}_S[||I-M||_F^2], \mathbf{E}_S[(x^T(I-M)x)^2], \mathbf{E}_S[x^T(I-M)^2x], \mathbf{E}_S[(x^T(I-M)x)\mathrm{Tr}(I-M)], \mathbf{E}_S[\mathrm{Tr}(I-M)^2]$
We first note that for for any r_1, r_2 and $a_1 \neq b_1, a_2 \neq b_2$ the quantity

$$\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}]$$

is only nonzero when $r_1 = r_2$ and $\{a_1, b_1, a_2, b_2\}$ contains two distinct elements, each with multiplicity 2 (let $\mathbf{I}_*(\{a_q, b_q\}_{q=1}^2)$ denote the indicator of the latter condition). In that case one has $\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}] = 1/B^2$. Note that the expression above appears in all of $\mathbf{E}_S[(x^T(I-M)x)^2], \mathbf{E}_S[x^T(I-M)^2x], \mathbf{E}_S[(x^T(I-M)x)\mathrm{Tr}(I-M)], \mathbf{E}_S[(\mathrm{Tr}(I-M))^2]$. Specifically, all of these expressions can be written as

$$\sum_{r_1=1}^{B}\sum_{r_2=1}^{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}]$$
$$\cdot (U_{a_1}U_{a_2}^T)^A(U_{b_1}U_{b_2}^T)^B\cdot((Ux)_{a_1}(Ux)_{a_2})^C((Ux)_{b_1}(Ux)_{b_2})^D\cdot((Ux)_{a_1}(Ux)_{b_1})^E(U_{a_1}U_{b_1}^T)^F\cdot((Ux)_{a_2}(Ux)_{b_2})^G(U_{a_2}U_{b_2}^T)^H,$$

where $A, B, C, D, E, F, G, H \in \{0, 1\}$ and $A + B + C + D + E + F + G + H = 2$. We thus have

$$|\sum_{r_1=1}^{B}\sum_{r_2=1}^{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}]\cdot$$
$$\cdot (U_{a_1}U_{a_2}^T)^A(U_{b_1}U_{b_2}^T)^B\cdot((Ux)_{a_1}(Ux)_{a_2})^C((Ux)_{b_1}(Ux)_{b_2})^D\cdot((Ux)_{a_1}(Ux)_{b_1})^E(U_{a_1}U_{b_1}^T)^F\cdot((Ux)_{a_2}(Ux)_{b_2})^G(U_{a_2}U_{b_2}^T)^H]|$$
$$\leq \frac{1}{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^2)|U_{a_1}U_{a_2}^T|^A|U_{b_1}U_{b_2}^T|^B\cdot|(Ux)_{a_1}(Ux)_{a_2}|^C|(Ux)_{b_1}(Ux)_{b_2}|^D\cdot$$
$$\cdot|(Ux)_{a_1}(Ux)_{b_1}|^E|U_{a_1}U_{b_1}^T|^F\cdot|(Ux)_{a_2}(Ux)_{b_2}|^G|U_{a_2}U_{b_2}^T|^H.$$

We have $|U_aU_b^T| \leq ||U_a||_2\cdot||U_b||_2$ by Cauchy-Schwarz, and $|(Ux)_a| \leq ||U_a||_2\cdot O(\sqrt{\log n})$ since $x \in \mathcal{T}^*$ by assumption of the lemma, so

$$\frac{1}{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^2)|U_{a_1}U_{a_2}^T|^A|U_{b_1}U_{b_2}^T|^B\cdot|(Ux)_{a_1}(Ux)_{a_2}|^C|(Ux)_{b_1}(Ux)_{b_2}|^D\cdot|(Ux)_{a_1}(Ux)_{b_1}|^E|U_{a_2}U_{b_2}^T|^F$$
$$\leq (O(\log n))^{C+D+E+G}\frac{1}{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^2)(||U_{a_1}||_2||U_{a_2}||_2)^A\cdot(||U_{b_1}||_2||U_{b_2}||_2)^B\cdot(||U_{a_1}||_2||U_{a_2}||_2)^C$$
$$\cdot(||U_{b_1}||_2||U_{b_2}||_2)^D\cdot(||U_{a_1}||_2||U||_{b_1}||_2)^E(||U_{a_1}||_2||U_{b_1}||_2)^F\cdot(||U_{a_2}||_2||U||_{b_2}||_2)^G(||U_{a_2}||_2||U_{b_2}||_2)^H.$$

Since we are only summing over $\{a_1, a_2, b_1, b_2\}$ that contain two distinct elements, we have

$$(O(\log n))^{C+D+E+G}\frac{1}{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^2)(||U_{a_1}||_2||U_{a_2}||_2)^A\cdot(||U_{b_1}||_2||U_{b_2}||_2)^B\cdot(||U_{a_1}||_2||U_{a_2}||_2)^C$$
$$\cdot(||U_{b_1}||_2||U_{b_2}||_2)^D\cdot(||U_{a_1}||_2||U||_{b_1}||_2)^E(||U_{a_1}||_2||U_{b_1}||_2)^F\cdot(||U_{a_2}||_2||U||_{b_2}||_2)^G(||U_{a_2}||_2||U_{b_2}||_2)^H$$
$$\leq (O(\log n))^{C+D+E+G}\frac{1}{B}\sum_{a_1,b_1=1}^{n}||U_{a_1}||_2^2||U_{a_2}||_2^2$$
$$\leq (O(\log n))^{C+D+E+G}\frac{1}{B}(\sum_{a_1=1}^{n}||U_{a_1}||_2^2)^2$$
$$\leq (O(\log n))^{C+D+E+G}\frac{d^2}{B},$$

where we used the fact that $\sum_a ||U_a|_2^2 = d$. Noting that $C+D+E+G = 0$ for $\mathbf{E}_S[||I-M||_F^2]$ and $C+D+E+G = 1$ for $\mathbf{E}_S[x^T(I-M)x\mathrm{Tr}(I-M)]$ completes the proof.

Bounding $\mathbf{E}_S[(x^T(I-M)x)^2\mathrm{Tr}(I-M)], \mathbf{E}_S[x^T(I-M)^2x\cdot\mathrm{Tr}(I-M)], \mathbf{E}_S[||I-M||_F^2\cdot\mathrm{Tr}(I-M)], \mathbf{E}_S[(x^T(I-M)x)^2\cdot x^T(I-M)x], \mathbf{E}_S[x^T(I-M)^2x\cdot x^T(I-M)x], \mathbf{E}_S[||I-M||_F^2\cdot x^T(I-M)x]$ All of the above expressions can be written as

$$\sum_{r_1=1}^{B}\sum_{r_2=1}^{B}\sum_{r_3=1}^{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}S_{r_3,a_3}S_{r_3,b_3}]$$
$$\cdot(U_{a_1}U_{a_2}^T)^A(U_{b_1}U_{b_2}^T)^B\cdot((Ux)_{a_1}(Ux)_{a_2})^C((Ux)_{b_1}(Ux)_{b_2})^D\cdot((Ux)_{a_1}(Ux)_{b_1})^E(U_{a_1}U_{b_1}^T)^F\cdot((Ux)_{a_2}(Ux)_{b_2})^G(U_{a_2}U_{b_2}^T)^H$$
$$\cdot((Ux)_{a_3}(Ux)_{b_3})^I(U_{a_3}U_{b_3}^T)^J$$

where $A, B, C\ldots$ are in $\{0,1\}$ and $A+B+C+D+E+F+G+H+I+J = 3$.

We first note that for for any r_1, r_2, r_3 and $a_1\neq b_1, a_2\neq b_2, a_3\neq b_3$ the quantity

$$\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}S_{r_3,a_3}S_{r_3,b_3}]$$

is only nonzero when $r_1 = r_2 = r_3$ and $\{a_1, b_1, a_2, b_2, a_3, b_3\}$ contains three distinct elements, each with multiplicity 2. Let $\mathbf{I}_*(\{a_q, b_q\}_{q=1}^3)$ denote the indicator of the latter condition. In that case one has $\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}S_{r_3,a_3}S_{r_3,b_3}] = 1/B^3$. Note we cannot have $a_1 = a_2 = a_3$ and $b_1 = b_2 = b_3$ since the expectation is 0 in that case.

Similarly to the above, it thus suffices to bound

$$\frac{1}{B^2}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^3)$$
$$\cdot|(U_{a_1}U_{a_2}^T)^A(U_{b_1}U_{b_2}^T)^B\cdot((Ux)_{a_1}(Ux)_{a_2})^C((Ux)_{b_1}(Ux)_{b_2})^D\cdot((Ux)_{a_1}(Ux)_{b_1})^E(U_{a_1}U_{b_1}^T)^F\cdot((Ux)_{a_2}(Ux)_{b_2})^G(U_{a_2}U_{b_2}^T)^H$$
$$\cdot((Ux)_{a_3}(Ux)_{b_3})^I(U_{a_3}U_{b_3}^T)^J|$$
$$\leq(O(\log n))^{C+D+E+G+I}\frac{1}{B^2}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^3)\cdot$$
$$\cdot(||U_{a_1}||_2||U_{a_2}||_2)^A(||U_{b_1}||_2||U_{b_2}||_2)^B\cdot(||U_{a_1}||_2||U_{a_2}||_2)^C(||U_{b_1}||_2||U_{b_2}||_2)^D\cdot(||U_{a_1}||_2||U_{b_1}||_2)^E(||U_{a_1}||_2||U_{b_1}||_2)^F$$
$$\cdot(||U_{a_2}||_2||U_{b_2}||_2)^G(||U_{a_2}||_2||U_{b_2}||_2)^H\cdot(||U_{a_3}||_2||U_{b_3}||_2)^I(||U_{a_3}||_2||U_{b_3}||_2)^J$$

where we used Cauchy-Schwarz and the assumption that $x\in\mathcal{T}^*$ (and hence x is not correlated with any of the rows of U too much), as above.

Since we are only summing over $\{a_1, a_2, a_3, b_1, b_2, b_3\}$ that contain three distinct elements, the expression above is upper bounded by

$$\begin{aligned}&(O(\log n))^{C+D+E+G+I}\frac{1}{B^2}\sum_{a,c,b}^{n}||U_a||_2^2||U_b||_2^2||U_c||_2^2\\ &\leq(O(\log n))^{C+D+E+G+I}\frac{d^3}{B^2}\\ &\leq(O(\log n))^2\frac{d^2}{B},\end{aligned}$$

where we used the fact that $\sum_a ||U_a|_2^2 = d$ and that in all cases, $C+D+E+G+I\leq 2$.

Bounding $\mathbf{E}_S[((x^T(I-M)x)^2 + x^T(I-M)^2x + ||I-M||_F^2 + (\mathrm{Tr}(I-M))^2)^2]$ and $\mathbf{E}_S[x^T(I-M)x \cdot \mathrm{Tr}(I-M)\cdot((x^T(I-M)x)^2 + x^T(I-M)^2x + ||I-M||_F^2 + (\mathrm{Tr}(I-M))^2)]$ All of the pairwise products arising in the expansion of the above expressions can be written as

$$\sum_{r_1=1}^{B}\sum_{r_2=1}^{B}\sum_{r_3=1}^{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\sum_{\substack{a_4,b_4=1,\\ a_4\neq b_4}}^{n}\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}S_{r_3,a_3}S_{r_3,b_3}S_{r_4,a_4}S_{r_4,b_4}]$$
$$\cdot (U_{a_1}U_{a_2}^T)^A(U_{b_1}U_{b_2}^T)^B\cdot((Ux)_{a_1}(Ux)_{a_2})^C((Ux)_{b_1}(Ux)_{b_2})^D\cdot((Ux)_{a_1}(Ux)_{b_1})^E(U_{a_1}U_{b_1}^T)^F\cdot((Ux)_{a_2}(Ux)_{b_2})^G(U_{a_2}U_{b_2}^T)^H$$
$$\cdot (U_{a_3}U_{a_4}^T)^{A'}(U_{b_3}U_{b_4}^T)^{B'}\cdot((Ux)_{a_3}(Ux)_{a_4})^{C'}((Ux)_{b_3}(Ux)_{b_4})^{D'}\cdot((Ux)_{a_3}(Ux)_{b_3})^{E'}(U_{a_3}U_{b_3}^T)^{F'}\cdot((Ux)_{a_4}(Ux)_{b_4})^{G'}(U_{a_4}U_{b_4}^T)^{H'},$$

where $A,B,C,D,E,F,G,H,A',B',C',D',E',F',G',H' \in \{0,1\}$ and add up to 4.

We now need to consider two cases.

Case 1: the number of distinct elements in $\{a_1,b_1,a_2,b_2,a_3,b_3,a_4,b_4\}$ is four, each occurring with multiplicity 2 (let $\mathbf{I}_*(\{a_q,b_q\}_{q=1}^4)$ denote the indicator of the latter condition) Then

$$\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}S_{r_3,a_3}S_{r_3,b_3}S_{r_4,a_4}S_{r_4,b_4}]$$

contributes $1/B^4$. In this case the number of distinct elements in $\{r_1,r_2,r_3,r_4\}$ cannot be larger than 2.

It thus suffices to bound

$$\frac{1}{B^2}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\sum_{\substack{a_4,b_4=1,\\ a_4\neq b_4}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^4)\cdot$$
$$\cdot |(U_{a_1}U_{a_2}^T)^A(U_{b_1}U_{b_2}^T)^B\cdot((Ux)_{a_1}(Ux)_{a_2})^C((Ux)_{b_1}(Ux)_{b_2})^D\cdot((Ux)_{a_1}(Ux)_{b_1})^E(U_{a_1}U_{b_1}^T)^F\cdot((Ux)_{a_2}(Ux)_{b_2})^G(U_{a_2}U_{b_2}^T)^H$$
$$\cdot (U_{a_3}U_{a_4}^T)^{A'}(U_{b_3}U_{b_4}^T)^{B'}\cdot((Ux)_{a_3}(Ux)_{a_4})^{C'}((Ux)_{b_3}(Ux)_{b_4})^{D'}\cdot((Ux)_{a_3}(Ux)_{b_3})^{E'}(U_{a_3}U_{b_3}^T)^{F'}\cdot((Ux)_{a_4}(Ux)_{b_4})^{G'}(U_{a_4}U_{b_4}^T)^{H'}|$$
$$\leq (O(\log n))^2\frac{1}{B^2}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^4)\cdot$$
$$\cdot (||U_{a_1}||_2||U_{a_2}||_2)^A(||U_{b_1}||_2||U_{b_2}||_2)^B\cdot(||U||_{a_1}||U_{a_2}||)^C(||U_{b_1}||_2||U_{b_2}||_2)^D\cdot(||U_{a_1}||_2||U_{b_1}||_2)^E(||U_{a_1}||_2||U_{b_1}||_2)^F$$
$$\cdot (||U_{a_2}||||U_{b_2}||)^G(||U_{a_2}||_2||U_{b_2}||_2)^H$$
$$\cdot (||U_{a_3}||_2||U_{a_4}||_2)^{A'}(||U_{b_3}||_2||U_{b_4}||_2)^{B'}\cdot(||U_{a_3}||_2||U_{a_4}||_2)^{C'}(||U_{b_3}||_2||U_{b_4}||_2)^{D'}\cdot(||U_{a_3}||_2||U_{b_3}||_2)^{E'}(||U_{a_3}||_2||U_{b_3}||_2)^{F'}$$
$$\cdot (||U_{a_4}||_2||U_{b_4}||_2)^{G'}(||U_{a_4}||_2||U_{b_4}||_2)^{H'}$$

where we used Cauchy-Schwarz and the assumption that $x \in \mathcal{T}^*$ (and hence x is not correlated with any of the rows of U too much), as above.

Since we are only summing over $\{a_1,a_2,a_3,a_4,b_1,b_2,b_3,b_4\}$ that contain three distinct elements, each of multiplicity two, the expression above is upper bounded by

$$(O(\log n))^2\frac{1}{B^2}\sum_{a,b,c,d}^{n}||U_a||_2^2||U_b||_2^2||U_c||_2^2||U_d||_2^2$$
$$\leq (O(\log n))^2\frac{d^4}{B^2}$$
$$\leq (O(\log n))^2\frac{d^2}{B},$$

where we used the fact that $\sum_a ||U_a|_2^2 = d$.

Case 2: the number of distinct elements in $\{a_1, b_1, a_2, b_2, a_3, b_3, a_4, b_4\}$ is two, each occurring with multiplicity 4 (let $\mathbf{I}_*(\{a_q, b_q\}_{q=1}^4)$ denote the indicator of the latter condition) Then

$$\mathbf{E}_S[S_{r_1,a_1}S_{r_1,b_1}S_{r_2,a_2}S_{r_2,b_2}S_{r_3,a_3}S_{r_3,b_3}S_{r_4,a_4}S_{r_4,b_4}]$$

contributes $1/B^2$. In this case the number of distinct elements in $\{r_1, r_2, r_3, r_4\}$ has to be one, since each column of S has a single non-zero entry and necessarily $a_1 = a_2 = a_3 = a_4$ and $b_1 = b_2 = b_3 = b_4$.

It thus suffices to bound

$$\frac{1}{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\sum_{\substack{a_4,b_4=1,\\ a_4\neq b_4}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^4)\cdot$$
$$\cdot|(U_{a_1}U_{a_2}^T)^A(U_{b_1}U_{b_2}^T)^B\cdot((Ux)_{a_1}(Ux)_{a_2})^C((Ux)_{b_1}(Ux)_{b_2})^D\cdot((Ux)_{a_1}(Ux)_{b_1})^E(U_{a_1}U_{b_1}^T)^F\cdot((Ux)_{a_2}(Ux)_{b_2})^G(U_{a_2}U_{b_2}^T)^H$$
$$\cdot(U_{a_3}U_{a_4}^T)^{A'}(U_{b_3}U_{b_4}^T)^{B'}\cdot((Ux)_{a_3}(Ux)_{a_4})^{C'}((Ux)_{b_3}(Ux)_{b_4})^{D'}\cdot((Ux)_{a_3}(Ux)_{b_3})^{E'}(U_{a_3}U_{b_3}^T)^{F'}\cdot((Ux)_{a_4}(Ux)_{b_4})^{G'}(U_{a_4}U_{b_4}^T)^{H'}|$$
$$\leq(O(\log n))^2\frac{1}{B}\sum_{\substack{a_1,b_1=1,\\ a_1\neq b_1}}^{n}\sum_{\substack{a_2,b_2=1,\\ a_2\neq b_2}}^{n}\sum_{\substack{a_3,b_3=1,\\ a_3\neq b_3}}^{n}\sum_{\substack{a_4,b_4=1,\\ a_4\neq b_4}}^{n}\mathbf{I}_*(\{a_q,b_q\}_{q=1}^4)\cdot$$
$$\cdot(||U_{a_1}||_2||U_{a_2}||_2)^A(||U_{b_1}||_2||U_{b_2}||_2)^B\cdot(||U||_{a_1}||U_{a_2}||)^C(||U_{b_1}||_2||U_{b_2}||_2)^D\cdot(||U_{a_1}||_2||U_{b_1}||_2)^E(||U_{a_1}||_2||U_{b_1}||_2)^F$$
$$\cdot(||U_{a_2}||||U_{b_2}||)^G(||U_{a_2}||_2||U_{b_2}||_2)^H$$
$$\cdot(||U_{a_3}||_2||U_{a_4}||_2)^{A'}(||U_{b_3}||_2||U_{b_4}||_2)^{B'}\cdot(||U_{a_3}||_2||U_{a_4}||_2)^{C'}(||U_{b_3}||_2||U_{b_4}||_2)^{D'}\cdot(||U_{a_3}||_2||U_{b_3}||_2)^{E'}(||U_{a_3}||_2||U_{b_3}||_2)^{F'}$$
$$\cdot(||U_{a_4}||_2||U_{b_4}||_2)^{G'}(||U_{a_4}||_2||U_{b_4}||_2)^{H'}$$

where we used Cauchy-Schwarz and the assumption that $x \in \mathcal{T}^*$ (and hence x is not correlated with any of the rows of U too much), as above.

Since we are only summing over $\{a_1, a_2, a_3, a_4, b_1, b_2, b_3, b_4\}$ that contain two distinct elements, each of multiplicity four, the expression above is upper bounded by

$$\begin{aligned}
&(O(\log n))^2\frac{1}{B}\sum_{a,b}^{n}||U_a||_2^4||U_b||_2^4\\
&=(O(\log n))^2\frac{1}{B}\sum_{a,b}^{n}||U_a||_2^2||U_b||_2^2 \qquad \text{(since } ||U_a||_2\leq 1 \text{ for all } a)\\
&\leq(O(\log n))^2\frac{d^2}{B},
\end{aligned}$$

where we used the fact that $\sum_a ||U_a|_2^2 = d$.

□

9 A counterexample for Fast Hadamard Transforms

A natural alternative transform to try would be the $m \times d$ Subsampled Randomized Hadamard Transform (SRHT) (see the references in Theorem 7 of Woodruff (2014)), which has the form $V = P \cdot H \cdot D$, where P is a diagonal matrix with a random subset of m diagonal entries equal to 1, and the remaining equal to 0, H is the Hadamard transform, and D is a diagonal matrix with random signs along the diagonal. Like FastFood, the SRHT can be applied to a d-dimensional vector in $O(d \log d)$ time. Note that each row of V is in the set $\{-1/\sqrt{d}, +1/\sqrt{d}\}$.

An illustrative counterexample would be to consider a pentagon inscribed in a unit circle 6 with one point p at $(0, 1)$. Each extreme point then receives $1/5$ of the circumference of the enclosing circle and so to be in the normal cone at p, one needs to have an angle in $[3\pi/10, 7\pi/10]$. Hence, the second coordinate ("y") needs to have magnitude at least $\sin(3\pi/10)$ which is larger than $1\sqrt{2}$ and so a vector in $\{-1\sqrt{2}, +1\sqrt{2}\}^2$ will never be in it. Generalizing

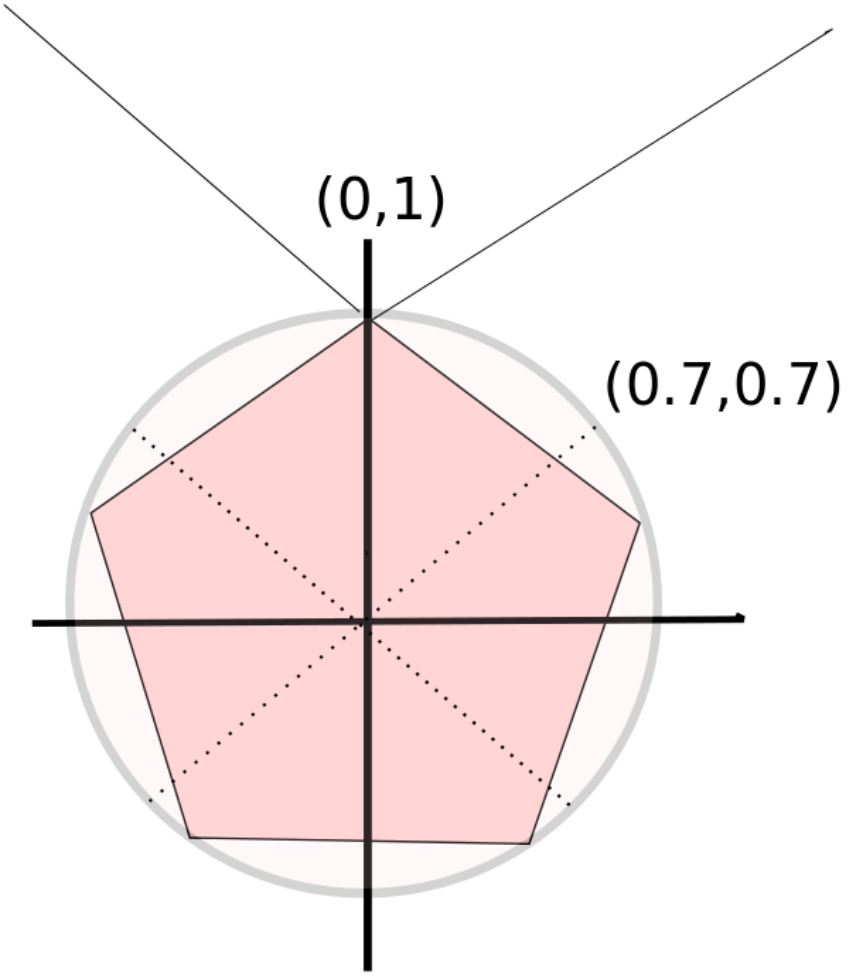

Figure 6: Illustrative example with 5 points corresponding the vertices of a pentagon inscribed in a unit circle. The normal cone at $(0, 1)$ is not covered by any of the vectors in $\{-1\sqrt{2}, +1\sqrt{2}\}^2$.

this to d-dimensions, we could consider a convex set C entirely supported on the first 2 coordinates (so 0 on the remaining coordinates). Further, we have that C is a pentagon with one extreme point equal to $(0, 1, \ldots, 0)$. Now we require the second coordinate to have magnitude at least $\sin(3\pi/10)$ which is larger than $1\sqrt{d}$ and therefore a vector in $\{-1\sqrt{d}, +1\sqrt{d}\}^d$ will never be in it. With probability ≈ 0.16, a random point on the sphere will have $x_1 > 1/\sqrt{d}$ (this corresponds to one standard deviation of an $N(0, 1/d)$ random variable), which means $\omega(N_C(p)) \approx 0.16$, yet *no row of* V will be in $N_C(p)$, which means that even if the condition number κ is constant, an algorithm using the SRHT in place of the FastFood transform will fail with probability 1.

10 SVM with Random Projections

We require the following stronger theorem for the SVM problem Paul et al. (2014).

Theorem 27. *Let $\epsilon \in (0, \frac{1}{2}]$ be an accuracy parameter and let $R \in \mathbb{R}^{d\times r}$ be a matrix satisfying $\|V^TV - V^TRR^TV\|_2 \leq \epsilon$ where $V \in \mathbb{R}^{d\times\rho}$ be the orthonormal (columns) matrix of right singular vectors obtained from the SVD of X. Let γ^**

and $\tilde{\gamma}^*$ be the margins obtained by solving the SVM problems using data matrices X and XR respectively. Then

$$(1-\epsilon)\gamma^{*2} \leq \tilde{\gamma}^{*2} \leq (1+\epsilon)\gamma^{*2}$$

Proof. We will follow a similar structure from Paul et al Paul et al. (2014). Let us denote by $E := V^TV - V^TRR^TV$ and for optimal solution vectors α^*, $\tilde{\alpha}^*$ the dual SVM objectives are given by:

$$Z_{\text{opt}} = 1^T\alpha^* - \frac{1}{2}\alpha^{*T}YXX^TY\alpha^* \tag{39}$$

$$\tilde{Z}_{\text{opt}} = 1^T\tilde{\alpha}^* - \frac{1}{2}\tilde{\alpha}^{*T}YXRR^TX^TY\tilde{\alpha}^* \tag{40}$$

Let us first consider the objective function at the optimal vector α of the original problem:

$$\begin{aligned} Z_{\text{opt}} &= 1^T\alpha^* - \frac{1}{2}\alpha^{*T}YXX^TY\alpha^* \\ &= 1^T\alpha^* - \frac{1}{2}\alpha^{*T}YU\Sigma V^TRR^TV\Sigma U^TY\alpha^* - \frac{1}{2}\alpha^{*T}YU\Sigma E\Sigma U^TY\alpha^* \\ &\geq \tilde{Z}_{\text{opt}} - \frac{1}{2}\tilde{\alpha}^{*T}YU\Sigma E\Sigma U^TY\tilde{\alpha}^* \end{aligned} \tag{41}$$

where we substituted the vector $\tilde{\alpha}^*$ in the objective and utilize the fact that it is smaller than the optimal value. Next, we consider the projected problem and lower-bound it as follows:

$$\begin{aligned} \tilde{Z}_{\text{opt}} &= 1^T\tilde{\alpha}^* - \frac{1}{2}\tilde{\alpha}^{*T}YXRR^TX^TY\tilde{\alpha}^* \\ &= 1^T\tilde{\alpha}^* - \frac{1}{2}\tilde{\alpha}^{*T}YU\Sigma V^TV\Sigma U^TY\tilde{\alpha}^* - \frac{1}{2}\tilde{\alpha}^{*T}YU\Sigma(-E)\Sigma U^TY\tilde{\alpha}^* \\ &\geq Z_{\text{opt}} - \frac{1}{2}\tilde{\alpha}^{*T}YU\Sigma(-E)\Sigma U^TY\tilde{\alpha}^* \end{aligned} \tag{42}$$

where we substituted the vector α^* in the objective as before and utilize the fact that it is smaller than the optimal value. By sub-multiplicativity, we have by using the fact $V^TV = I$:

$$\frac{1}{2}z^TYU\Sigma G\Sigma U^TYz \leq \frac{1}{2}\|z^TYU\Sigma\| \cdot \|G\|_2 \cdot \|\Sigma U^TYz\| \tag{43}$$

$$= \frac{1}{2}\|G\|_2 \cdot \|z^TYX\|_2^2 \tag{44}$$

for any vector z and matrix G. Let us bound the following second-order term:

$$\begin{aligned} |z^TYXRR^TX^TYz - z^TYXX^TYz| &= |z^TYU\Sigma(V^TRR^TV - V^TV)\Sigma U^TYz| \\ &= |z^TYU\Sigma(-E)\Sigma U^TYz| \\ &\leq \|E\|_2 \cdot \|z^TYU\Sigma\|_2^2 \\ &= \|E\|_2 \cdot \|z^TYX\|_2^2 \end{aligned}$$

for any vector z. This gives us the following useful inequality:

$$\|z^TYX\|_2^2 \leq \frac{1}{1-\|E\|_2}\|z^TYXR\|_2^2 \tag{45}$$

Combining (41), (42), (43), and (45) for $z \in (\alpha, \tilde{\alpha})$ and $G \in (-E, E)$, we have the following bounds:

$$\begin{aligned}
\tilde{Z}_{\text{opt}} &\geq Z_{\text{opt}} - \frac{1}{2}\|E\|_2 \cdot \|\alpha^{*T} Y X\|_2^2 \\
&= Z_{\text{opt}} - \|E\|_2 \cdot Z_{\text{opt}} \\
&= (1 - \|E\|_2) Z_{\text{opt}} \qquad (46) \\
Z_{\text{opt}} &\geq \tilde{Z}_{\text{opt}} - \frac{1}{2}\|E\|_2 \cdot \|\tilde{\alpha}^{*T} Y X\|_2^2 \\
&\geq \tilde{Z}_{\text{opt}} - \frac{\|E\|_2}{1 - \|E\|_2} \frac{\|\tilde{\alpha}^{*T} Y X R\|_2^2}{2} \\
&= \tilde{Z}_{\text{opt}} - \frac{\|E\|_2}{1 - \|E\|_2} \tilde{Z}_{\text{opt}} \\
&= (1 - \frac{\|E\|_2}{1 - \|E\|_2}) \tilde{Z}_{\text{opt}} \qquad (47)
\end{aligned}$$

The bounds follow by using the following relations, $Z_{\text{opt}} = \frac{1}{2\gamma^{*2}}$ and $\tilde{Z}_{\text{opt}} = \frac{1}{2\tilde{\gamma}^{*2}}$:

$$(1 - \frac{\|E\|}{1 - \|E\|_2})\gamma^{*2} \leq \tilde{\gamma}^{*2} \leq \frac{1}{1 - \|E\|_2}\gamma^{*2} \qquad (48)$$

Notice that we cannot now trivially project the data to zero (XR to the zero matrix) which would have been acceptable if we had the used the weaker version of theorem as stated in Paul et al. (2014). □

11 Discussion

We have presented an efficient way to multiply by a Gaussian matrix, without actually computing the dense matrix product. Theorem 1 provides our theoretical guarantees on this much faster transform, showing it has low variation distance to multiplication by a dense Gaussian matrix.

Our transform is useful in a surprising number of applications — here we apply our transform to NMF and SVM. The classical way of speeding up Gaussian transforms via the Fast Hadamard or FFT does not work in our setting since it misses large sections of the sphere.

Our experiments on synthetic and real-world datasets for NMF showed that the results obtained by our algorithm were on par with the state-of-the-art NMF algorithms such as SC, XRAY and SPA. In particular, for synthetic problems, we showed similar anchor recovery performance as random projection (GP) of Damle and Sun Damle and Sun (2014) both in the noiseless and noisy cases. Also, the performance was remarkably similar to GP when applied on the breast cancer dataset and also picked up activation patterns which might be of biological interest as previously noted in flow cytometry problems. Experiments on document classification tasks using the popular SVM formulation revealed that the new projection leads to faster SVM solutions than previous methods. Previously, it was shown that while CountSketch led to faster projection times it did not lead to overall faster training time and was in fact was found to be slower than random Gaussian projections (RG). Our new countGauss projection fixes this by sacrificing projection time compared to countSketch projection but leads to an overall faster SVM training time and thereby beats both random Gaussian and CountSketch-based SVM algorithms Paul et al. (2014). We note that in practice for SVM, solution accuracy may be of critical importance rather than computation time and in these scenarios random projection based algorithms can be used to explore the optimal settings of the SVM parameters such as soft-margin. In our experiments (not shown) we noticed that these lead to faster training times while not sacrificing test accuracy.

Acknowledgements

Vamsi P. would like to acknowledge support from the following grant: NSF-IIP-1346452 and also RDI^2 at Rutgers University where this work was initiated. David W. would like to acknowledge the support from XDATA program

of the Defense Advanced Research Projects Agency (DARPA), administered through Air Force Research Laboratory contract FA8750-12-C-0323.

References

M. C. U. Araújo, T. C. B. Saldanha, R. K. H. Galvão, T. Yoneyama, H. C. Chame, and V. Visani. The successive projections algorithm for variable selection in spectroscopic multicomponent analysis. *Chemometrics and Intelligent Laboratory Systems*, 57(2):65–73, 2001.

S. Arora, R. Ge, R. Kannan, and A. Moitra. Computing a nonnegative matrix factorization–provably. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 145–162. ACM, 2012.

H. Avron, P. Maymounkov, and S. Toledo. Blendenpik: Supercharging lapack's least-squares solver. *SIAM J. Scientific Computing*, 32(3):1217–1236, 2010.

A. R. Benson, J. D. Lee, B. Rajwa, and D. F. Gleich. Scalable methods for nonnegative matrix factorizations of near-separable tall-and-skinny matrices. In *Advances in Neural Information Processing Systems*, pages 945–953, 2014.

G. Brumfiel. Down the petabyte highway. *Nature*, 469(20):282–283, 2011.

K. L. Clarkson and D. P. Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 81–90, New York, NY, USA, 2013a. ACM. ISBN 978-1-4503-2029-0. doi: 10.1145/2488608.2488620. URL http://doi.acm.org/10.1145/2488608.2488620.

K. L. Clarkson and D. P. Woodruff. Low rank approximation and regression in input sparsity time. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 81–90, 2013b.

D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comput.*, 9(3):251–280, 1990.

A. Damle and Y. Sun. Random projections for non-negative matrix factorization. *arXiv preprint arXiv:1405.4275*, 2014.

C. Ding, T. Li, W. Peng, and H. Park. Orthogonal nonnegative matrix t-factorizations for clustering. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 126–135. ACM, 2006.

D. Donoho and V. Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in neural information processing systems*, page None, 2003.

N. Gillis and F. Glineur. Using underapproximations for sparse nonnegative matrix factorization. *Pattern recognition*, 43(4):1676–1687, 2010.

N. Gillis, S. Vavasis, et al. Fast and robust recursive algorithmsfor separable nonnegative matrix factorization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 36(4):698–714, 2014.

A. Kumar, V. Sindhwani, and P. Kambadur. Fast conical hull algorithms for near-separable non-negative matrix factorization. In *Proceedings of The 30th International Conference on Machine Learning*, pages 231–239, 2013.

Q. Le, T. Sarlós, and A. Smola. Fastfood–approximating kernel expansions in loglinear time. In *Proceedings of the international conference on machine learning*, 2013.

D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pages 556–562, 2001.

P. Paatero and U. Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2):111–126, 1994.

S. Paul, C. Boutsidis, M. Magdon-Ismail, and P. Drineas. Random projections for linear support vector machines. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 8(4):22, 2014.

B. Recht, C. Re, J. Tropp, and V. Bittorf. Factoring nonnegative matrices with linear programs. In *Advances in Neural Information Processing Systems*, pages 1214–1222, 2012.

M. Tepper and G. Sapiro. Compressed nonnegative matrix factorization is fast and accurate. *arXiv preprint arXiv:1505.04650*, 2015.

J. A. Tropp. Improved analysis of the subsampled randomized hadamard transform. *Advances in Adaptive Data Analysis*, 3(1-2):115–126, 2011.

S. A. Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2009.

V. V. Williams. Multiplying matrices faster than coppersmith-winograd. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 887–898, 2012.

D. P. Woodruff. Sketching as a tool for numerical linear algebra. *arXiv preprint arXiv:1411.4357*, 2014.

K. Y. Yılmaz, A. T. Cemgil, and U. Simsekli. Generalised coupled tensor factorisation. In *Advances in Neural Information Processing Systems*, pages 2151–2159, 2011.